University of Alberta

Prototyping Human Perception-Action Systems

by

Gordon J Binsted

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Department of Psychology

Edmonton, Alberta

Fall, 2001

National Library of Canada
Acquisitions and Bibliographic Services
395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque nationale du Canada
Acquisitions et services bibliographiques
395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-68911-5

Canada

University of Alberta

Library Release Form

Name of Author: *Gordon James Binsted*

Title of Thesis: *Prototyping Human Action-Perception Systems*

Degree: *Doctor of Philosophy*

Year this Degree Granted: *2001*

Permission is hereby granted to the University of Alberta Library to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purposes only.

The author reserves all other publication and other rights in association with the copyright in the thesis, and except as herein before provided, neither the thesis nor any substantial portion thereof may be printed or otherwise reproduced in any material form whatever without the author's prior written permission.

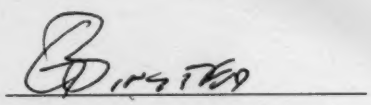

10945 74th Avenue
Edmonton, AB
Canada, T6G 0E5

Aug. 14, 2001

University of Alberta

Faculty of Graduate Studies and Research

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled *Prototyping Human Perception-Action Systems* submitted by *Gordon James Binsted* in partial fulfillment of the requirements for the degree of *Doctor of Philosophy*.

Dr. Terry Caelli

Dr. Romeo Chua

Dr. Walter Bischof

Dr. Brian Maraj

Dr. Pete Dixon

Dr. Steve Wallace

Dr. Martin Jägersand

Aug 13, 2001

0.1 Abstract

Computational models of motor control vary widely in method and application, often demanding extensive calculations in order to minimize cost functions and/or extensive implicit knowledge of physical properties of the entire system (e.g., Kawato et al, 1992, 1996). Although such models have had demonstrated success in predicting wide ranges of movement behaviour they are limited in their ability to capture the variability of performance normally displayed by the human system. Current models are similarly unable to account for more than a confined task type, often making restrictions such as limiting performance to open-loop control of simple reaching movements. Presented here are applications of hidden Markov models (HMMs) - a tool known in the machine learning literature for representing first-order dynamical systems in a stochastic fashion - and its generalization, dynamical Bayesian networks (DBN), to human movement. The HMM approach enables the representation of probabilistic relations between elements of a system for the expression of system dynamics; the structure of DBNs is such that combinations of HMMs may be used to build a model of the target system based on biological knowledge, or hypothesis. An extensive discussion of methods for training and assessing both HMMs and DBN is presented within a motor control context. Further, four experiments are described which examine the utility and efficacy of these methods for the representation, recognition, and production of both discrete and continuous motor tasks. General comparisons are made throughout between DBN/HMM techniques and other modeling alternatives for representing

time-varying biological movement signals. Extensions of the HMM/DBN framework are provided along with suggestions for future applications and possible implications to current theories of motor control.

0.2 On How to Reach for a Cup of Coffee

Consider for a moment how to reach out for the cup of coffee on your desk. One way to approach this problem is to direct another person, who is blindfolded and unaware of the goal, to complete the task. You will likely use commands such as *forward*, *up*, *close your hand*; however, attempt the task by commanding each joint (use only *flex* and *extend*). While this task is far less trivial, the problem would seem almost impossible if one were to have to command each muscle (fibre?) with regard to information such as joint viscosity, muscle insertion position and angle, tissue length and velocity.

The ease with which humans move about and interact with their environment has been a compelling topic of research for over a century (e.g., Woodworth 1899). The simple acts of reaching toward an object or walking across a room require the complex integration of perceptual (e.g., visual and kinesthetic) and motor (e.g., kinetic and dynamic) processes. How are action and perception coordinated with such ease? Even more perplexing is that such performance occurs in an environment defined by complex physical principles which are not known, at least at a conscious level, by most people (or animals). How does the motor system account for such properties as gravity and friction, which mankind has only recently come to understand (on a evolutionary time-scale)?

Many researchers have examined exactly this question, formulating explanations from a variety of theoretical and methodological standpoints. However, a conspicuous feature

of much of this research is the incompatibility of the experimental results – often referred to as the "levels of analysis problem" – whereby hypotheses from one viewpoint are untestable or philosophically irrelevant to the other. The work presented in this dissertation addresses exactly this problem. The focus here is on developing a framework for prototyping motor behaviour and control that is accessible to a wide variety of scientists. In order to adequately develop this approach, a representative review of current models is supplied, giving a reference for the strengths and weaknesses of extant methods in motor control. Similarly, the first chapter discusses some of the seminal research in motor control to highlight critical issues on control and problems that must be addressable by any general model.

The proposed framework is based on methods of Bayesian network modeling. This approach originates in the machine learning / pattern recognition literature where implementations of hidden Markov models (HMMs), a subset of Bayesian networks, have been applied to modeling a wide range of spatio-temporal data including handwriting (e.g., Marti & Bunke 2001, Park, Sin, Moon & Lee 2001), gesture recognition (e.g., Wilson & Bobick 2001) and surveillance (Bui, Venkatesh & West 2001). While this tool has been used widely, its applicability to problems of biological *control* has received little or no investigation. Further, the utility of coupling sets of HMMs (dynamical Bayesian networks) to prototype biological dynamical systems has not been examined quantitatively.

Although extensive discussions throughout this dissertation centre on the biological

implications of the presented Bayesian models, the goal is to assess the viability of Bayesian modeling *as a technique* for encapsulating key features of human performance. Even if the assumptions of this approach should prove biologically untenable as a process, HMMs/DBNs are a statistical representation of the target system. This being the case, the internal structure of the models encapsulate central tendencies of performance ($\bar{x}$ and σ) for experimentalists wishing to understand the impact of their independent manipulations, or for clinicians wishing to evaluate their patients.

0.3 Acknowledgements

I would like to express my warmest thanks to my supervisors, Dr. Romeo Chua and Dr. Terry Caelli.

Special thanks also to my *unofficial* supervisors (and committee members), Dr. Walter Bischof and Dr. Brian Maraj, for their patience and reassuring words on a day-to-day basis.

I also gratefully acknowledge the support of the Department of Psychology and the Faculty of Physical Education at the University of Alberta. Without whose help I may never have found a home (department).

Last, but certainly not least, thanks to Gillian. A great wife and friend who has followed me around the continent in search of these three little letters. I would not have made it here without you.

The research described in this dissertation was completed with the support of the Natural Sciences and Engineering Research Council of Canada and the Isaac Walton Killam Memorial Trust.

Contents

0.1 Abstract i

0.2 On How to Reach for a Cup of Coffee iii

0.3 Acknowledgements vi

1 Perception and Action **1**

1.1 A Brief History of Perception-Action 2

1.1.1 A New Theory of Information Pickup 2

1.1.2 An Interdisciplinary Approach to Action 3

1.1.2.1 The Degrees of Freedom Problem 3

1.1.2.2 The Principle of Equal Simplicity 4

1.2 Approaches to the Study of Movement 5

1.2.1 The Information Processing Approach 5

1.2.1.1 Stimulus Encoding 6

1.2.1.2 Response Selection 6

1.2.1.3 Response Programming 7

1.2.1.4 Response Execution 8

1.2.1.5 The Role of Feedback 9

1.3 Gibsonian Approaches: Branches of Ecological Psychology 10

1.3.1 Direct Perception 10

1.3.2 Kinetic Theory 11

1.3.3 Pattern Dynamics (Dynamical Systems Approach) 12

2 Current Models in Motor Control **14**

2.1 The Optimized Sub-Movement Models 15

2.2 The Knowledge Model 18

2.2.1 Storage 19

2.2.2 Planning 19

2.2.3 Execution 21

2.3 Minimum Torque-Change Model 23

2.4 Equilibrium-Point Models 26

2.4.1 The γ- and α-Models 27

2.4.2 The λ-Model 28

2.5 A Kinematic Theory of Rapid Human Movement 30

2.6 The *VITE* Model 33

2.7 The Haken-Kelso-Bunz Model 35

2.8 Conclusions 40

2.9 Summary 42

3 A Computational Intelligence Perspective **44**

3.1 Basic Methods in Artificial Neural Networks 46

3.1.1 Perceptrons 47

3.1.2 Backpropagation 48

3.1.3 Recurrent Networks 50

3.1.4 Applications in Motor Control 51
3.1.4.1 *VITE* Revisited 51
3.1.4.2 Minimization Procedures 54
3.1.4.3 Neural Networks - Parting Thoughts 55
3.2 First-Order Markov Chains 56
3.3 Hidden Markov Models 57
3.3.1 Definitions 58
3.3.2 Classical Estimation Problems in HMMs 59
3.3.3 The Viterbi and Baum-Welch Procedures 60
3.3.3.1 The Forward Operator 61
3.3.3.2 The Backward Operator 62
3.3.3.3 The Viterbi Algorithm 63
3.3.3.4 The Baum-Welch Algorithm 65
3.3.4 Assessing HMM Performance 67
3.3.4.1 Unconstrained Monte Carlo Sampling 67
3.3.4.2 Constrained MonteCarlo Sampling 68
3.3.4.3 Hamming Distance 68
3.3.5 Hidden Markov Models – Closing Thoughts 69

4 Hidden Markov Models in Action **70**

4.1 Experiment 1: The Simplest Case 70
4.1.0.1 Hidden Markov Models for One-dimensional Aiming . . 70
4.1.1 Methods 73
4.1.1.1 Participants 73

4.1.1.2 Task and Procedure 73
4.1.1.3 Apparatus 73
4.1.1.4 Model Construction 74
4.1.1.5 Model Assessment 74
4.1.2 Results 75
4.2 Discussion 78
4.3 Toward a More Generalizeable Framework 80

5 Methods for Increasing HMM Utility **82**

5.1 Bayesian Networks 83
5.2 Dynamical Bayesian Networks 84
5.2.1 gDBN Structure 85
5.2.1.1 Generalized **A** Matrix 86
5.2.1.2 Generalized Observations (**B**) 87
5.2.1.3 Weighting Matrix 87
5.2.2 Estimate Update and Prediction 89
5.2.2.1 The Generalized Baum Welch Algorithm 89
5.2.2.2 Model Re-estimation 90
5.2.3 Assessing gDBN Performance 92
5.2.3.1 Prediction: The generalized Viterbi algorithm (gViterbi). 92
5.2.3.2 Monte Carlo Sampling of gDBNs 94
5.3 Other Approaches to gDBNs 96
5.3.1 N-Heads Approach 96
5.3.2 Residual Approach 98

5.3.3 gDBNs for Action 98

6 Dynamical Bayesian Networks in Action **100**

6.1 Experiment 2: A Time-Indexed Reach 100

6.1.1 Introduction 100

6.1.1.1 A Metronome for Action 101

6.1.1.2 A gDBN for Reaching 102

6.1.2 Methods 102

6.1.2.1 Participants 102

6.1.2.2 Task and Procedure 103

6.1.2.3 Apparatus 103

6.1.2.4 Model Construction 103

6.1.2.5 Model Assessment 104

6.1.3 Discussion and Results 106

6.1.3.1 Model Generalization 109

6.1.4 Conclusions 111

6.2 Experiment 3: Hypothetical Coupled Oscillators 112

6.2.1 Introduction 112

6.2.2 Model Construction 113

6.2.2.1 Deterministic Oscillators 115

6.2.2.2 Semi-deterministic Oscillators 115

6.2.2.3 Noisy Oscillators 115

6.2.3 Model Performance 116

6.2.4 Conclusions 122

7 On Learning the Internal Structure of a Biological Oscillator **124**

7.1 Experiment 4: Modeling a Human Oscillator system 124

7.1.1 Methods 126

7.1.1.1 Participants 126

7.1.1.2 Task 126

7.1.1.3 Apparatus 127

7.1.1.4 Procedure 127

7.1.1.5 Continuous Phase Angle & Relative Phase Estimates . 127

7.1.1.6 Data Summary and Clustering 128

7.1.1.7 Model Construction and Assessment 129

7.1.2 Discussion and Results 130

7.1.3 Conclusions 135

8 On Reaching for a Cup of Coffee: Revisited **138**

8.1 Extensions 139

8.1.1 Improving Ω Through gBaum Welch 139

8.1.2 A 'Local' gViterbi Controller 141

8.2 Concluding Remarks 142

A Polhemus Operations **157**

A.1 System Configuration 157

A.2 System Calibration 158

A.2.1 Gaussian Mixtures 159

A.3 A Neural Network Calibration Algorithm 160

List of Figures

1.1 Movement Parameter Scaling 8

2.1 Visual feedback models 17

2.2 The Knowledge Model 22

2.3 An Isolated Muscle-Spring 27

2.4 The Λ-Model 29

2.5 The Kinematic Theory 31

2.6 The *VITE* circuit 34

2.7 Relative Phasing of Fingers 36

2.8 The HKB-Model: Representations of model parameters 37

3.1 Basic Perceptron 47

3.2 Multi-Layer Neural Network 49

3.3 A Recurrent Neural Network Unfolded 52

3.4 The AVITE Circuit 53

3.5 Forward Procedure 62

3.6 Backward Procedure 63

3.7 Viterbi Search 64

3.8 Baum-Welch Algorithm 66

4.1 Elbow Rotation Task 72

4.2 Joint Angle Definitions 74

4.3 Unconstrained MonteCarlo (78^o) motion 76

4.4 Comparison: Simulated versus Recorded 77

5.1 Automotive Baysian Network 84

5.2 Oscillator DBN 85

5.3 Graphical 3-node DBN 88

5.4 Matrix 3-node DBN 89

5.5 Mass Action versu N-Heads for a DBN 97

6.1 gDBN Simulation: 100% Coupling 108

6.2 gDBN Simulation: 10% Coupling 108

6.3 gDBN Simulation: 0.1% Coupling 109

6.4 Amplitude Scaling Simulation: 10^o 109

6.5 Amplitude Scaling Simulation: 20^o 110

6.6 Temporal Scaling Simulation 111

6.7 Graphical Oscillator Model 114

6.8 Deterministic Oscillator DBN Performance 119

6.9 Stochastic Oscillator gDBN Performance (1%,0.1%) 120

6.10 Stochastic Oscillator gDBN Performance (10%,10%) 121

7.1 Discrete Observation Clustering 129

7.2 Experiment 4: Example Phase Transitions 133

A.1 The Gaussian Field Mixture Model 160

A.2 Neural network calibration model 161

List of Tables

2.1 Summary of current models in motor control 41

4.1 Experiment 1 Summary 81

6.1 Experiment 2 Participant Summary 107

6.2 Experiment 2 Model Summary 107

6.3 Deterministic Oscillator Structure 116

6.4 Stochastic Oscillator Structure 117

6.5 Noisy Oscillator Structure 118

7.1 Experiment 4: Summary Participant Performance 131

7.2 Anti-Phase Model 132

7.3 Transition Model 134

7.4 In-Phase (post-transition) Model 135

7.5 In-Phase Model 135

Chapter 1

Perception and Action

Historically, a majority of studies of human perception and action have considered these processes as separate and largely independent. However, a parallel viewpoint stresses perception-action coupling, where the interdependence and interplay between the sensory environment and the motor system of animal (human) is the focus. Over the last 50 years, two implementations of this perception-action approach have revolutionized the study of human movement. One, proposed by J.J. Gibson, approached perception with the goal of removing the need for complex internal representations within an animal. Instead, a theory of perception was presented that stressed the interplay between animal and environment and the active nature of perception. The other, formulated by N. Bernstein, posed the analysis of movement by presenting problems, and solutions, based on an interdisciplinary method – seeking complete understanding of behaviour in the combined approaches of Biology, Physics, and Psychology.

1.1 A Brief History of Perception-Action

1.1.1 A New Theory of Information Pickup

The primary thesis of Gibson (Gibson 1966, Gibson 1979) was that traditional theories of perception need to be abandoned in favour of an approach that reduces the need for either substantial cognitive preprocessing or meanings supplied by past experience or inheritance. More specifically, Gibson presents the idea that information may be picked up by the senses in the *active* process of perceiving – picking up information by exploring tiny fluxes of energy at receptors.

Gibson also provided new categories for *what* is perceived. Instead of breaking the world into colour, form, and motion, for example, the environment may be divided into groups based upon what they afford the animal and its intent (i.e., actions/properties which are possible given the structure):

> **Places:** A location or series of locations that make up the environment of the perceiver. A *place* affords no natural boundaries. Similarly, places are immobile, differentiating them from objects. Places primarily afford the act of locomotion.
>
> **Objects:**
>
> *Attached Objects:* An attached object is connected to a place, therefore it may not be moved.
>
> *Detached Objects:* A detached object is fully independent of place and thus may move; however, the common features of an object (e.g. colour and boundaries) remain.

While these categories may seem trivial divergences from the traditional viewpoint, they allow properties to be attached to objects based upon the group they belong to. For example, a book is a detached object; thus we may attribute the ability to

retain properties such as colour and texture, independent of place as well as affording movement (in this case, perhaps grasping). In order to *perceive* the environment, an animal detects both contiguity and variation in places and objects. Specifically, "the perceiver extracts the invariants of structure from the flux of stimulation while still noticing the flux" (Gibson 1979, p. 247).

By acknowledging the importance of environmental features and their impact on perception, Gibson provides a basis for the study of behaviour. In order to understand behaviour, researchers may investigate the physical properties of the system to extract the variations or consistencies which result in the observed behaviour. By placing perceptual information in the environment, the *subjective* task of perception becomes a predictable result of a objective event – instead of being due to an ambiguous internal representation of reality.

1.1.2 An Interdisciplinary Approach to Action

The Russian physiologist Nicholai Bernstein produced what has now become a landmark book in the areas of applied physiology and kinesiology, focusing on the study of human movement (Bernstein 1967). While Bernstein came from a physiology background, his approach to human movement stressed the important contributions of many disciplines including psychology, mathematics and physics.

1.1.2.1 The Degrees of Freedom Problem

A central idea discussed by Bernstein was the "degrees of freedom problem". This concept can most easily be demonstrated by discussing the control of a simple motor task, e.g., pointing. Consider a simple motor system composed of four linked rigid bodies (torso, upper-arm, lower-arm, hand) with joints containing two, one, and two degrees of freedom, respectively. There are five mechanical degrees of freedom (dfs) for

the motor system to control, whereas an object only has three dfs, x,y, and z. This example is dramatically over-simplified, the actual number of dfs for any movement actually includes the joints of the entire body (in the order of 100 dfs).* The excess in dfs results in the movement problem or task having infinite possible solutions. How does the motor system select from the alternatives to generate apparently orderly movement?

1.1.2.2 The Principle of Equal Simplicity

A second central idea presented by Bernstein addresses the questions of what is really represented in the brain and what is an appropriate taxonomy of actions (i.e., what actions are really similar/different?). Bernstein suggests that a line of *equal simplicity* is observed when the transition between tasks generates no change in movement accuracy, rate, or variability (i.e, simplicity). For example, it is equally simple for someone to use the dominant hand to write their name at many scales, even though these tasks could require the contributions of substantially different musculature. However, if the foot were used, the accuracy and variability of performance changes substantially.

If the idea of "equal simplicity" were incorrect, the corollary would be that specific muscles are controlled by specific regions of the brain. Thus, acquired movements would be difficult to transfer between groups of muscles. As this extension is demonstrably untrue†, Bernstein surmised that movements should be classified independently of muscles or muscle commands. Although this final point seems redundant within the literature of today where discussions of ideas like *generalized motor programs* (Schmidt 1982) are commonplace, it was presented at a time where localizing function was paramount to motor control research.

---

*The degrees of freedom problem may even be extended further if the additional df involved with muscle selection are considered.

†Examination of an introductory physiology text quickly dispells the idea of direct connections between specific cortical units and specific muscles, even though somatotopic representations certainly exist. Similarly, the learning literature stresses the *ability* of the motor system to transfer learning between muscles

1.2 Approaches to the Study of Movement

In order to systematically address problems of perception and action, the research community has adopted two views. One describes the flow of information through stages, with descriptions remaining relatively independent of physiological mechanisms[‡]. The second view follows approaches of Gibson and Bernstein. Here, events are modeled as occuring in a self-organising fashion (i.e, as a result the interaction between elements of the system), removing the need for an *informed* homunculus to manage internal events. However, in order for any approach to be accepted as a complete account of movement control it must accurately predict the planning and execution of actions as well as the manner by which information about the movement and the environment is utilized.

1.2.1 The Information Processing Approach

The information processing (IP) approach has a long history in the psychological literature, basing its results primarily upon stimulus-response paradigms. Although not initially intended for application beyond the initial reaction to the stimulus, IP has made the transition to full movement responses with relative ease. The primary measure, reaction time, is based on the assumption that processing proceeds in a primarily serial[§] manner with the corollary expectation that all processing events require time proportional to load and/or difficulty. Allowing for these assumptions, any task may be divided into discrete events: in movement these are stimulus encoding, response selection, response programming and response execution. However, the reliance upon elapsed time to delineate events represents a significant shortfall of all IP based theories; events may occur without temporally identifiable features.

[‡]Although most researchers do not address mechanism directly, imaging and lesioning techniques are often used to place a *locus* on events.

[§]This serial assumption is not necessarily stictly adhered to. Some IP views permit parallel streams of processing to occur until a processing bottleneck is reached (e.g., following stimulus encoding).

1.2.1.1 Stimulus Encoding

Stimulus encoding, the initial stage of any response, is sometimes subdivided into stimulus detection and stimulus identification. These sub-stages represent the processes involved with the transformation of the physical stimulus into an internal (biological) representation. Some important factors influencing processing time include stimulus intensity and clarity. While there are many implications of this stage to movement, the majority of principles are common with traditional *perceptual* IP.

1.2.1.2 Response Selection

Upon detection of the relevant stimuli, the responder must select the correct response from a number of possible ones – a choice reaction paradigm. Perhaps the most famous account of this behaviour comes in the form of Hick's Law (Hick 1952, Hyman 1953, see Equation 1.1) which suggests a linear relation between forced choice reaction time and the number of response alternatives.

$$ChoiceRT = a + b[log_2(N)] \tag{1.1}$$

Information processing researchers reduce this relation to an expression of the maximum *rate* of information processing (Equation 1.2) , where the amount of information is determined by (a) the amount of uncertainty in the environment prior to the stimulus, and (b) the reduction in uncertainty given the stimulus occurrence (Schmidt 1988). In this sense, information H may be expressed as:

$$H = log_2 \frac{1}{P_i} \tag{1.2}$$

where P_i is the probability of an event (i) occurring. Based on these simple relations, a stimulus environment may be expressed as containing a finite number of *bits* of information relating to the number of choices required to define a response. Thus, in

Hick's Law, the a term may be explicitly defined as the simple reaction time, and the $[Log_2(N)]$ term defined as the amount of information being received by the performer. This leaves b to represent the rate of information passage. Therefore, assuming a constant for b, as the environment presents a more variable informational array the time occupied by the selection process will increase.

1.2.1.3 Response Programming

Once the appropriate action has been selected, it must be planned in such a manner as to succeed in the task. Following the assumption of serial processing that carries throughout IP theories, there is an expectation for programming time to increase as the amount of processing increases (Henry & Rogers 1960). Also, while the process of movement programming seems conceptually valid, any explanation of the process must deal with two key issues. First, the storage problem suggests that maintaining a *program* for every action would require the motor system to store a very large number of instructions to allow movements to occur 'on command'. Second, the novelty problem poses how a program based system presumably based upon memory of movements, would generate previously unattempted actions (Schmidt 1988).

The most complete account of motor programs comes in the form of *Generalized Motor Programs* (GMP) (Schmidt 1976, Schmidt 1985, Schmidt 1988). A GMP is a set of *a priori* defined instructions for the completion of a *class* of actions. As such, there are features invariant to all movements served by a GMP as well as aspects (parameters) which are free to vary depending upon the instance of the action. Specifically, event order (e.g., muscle activation) is thought to remain constant across a class of movements specified by a GMP (see Figure 1.1). Similarly, the relative timing and relative magnitudes of events are thought to be invariant (e.g., duration of EMG activity). Some parameters are however free to vary; the overall duration of the movement, the absolute level of muscle activation/force output, and even the effector may be specified

within a program. An often cited demonstration of motor programs is human handwriting (Schmidt 1988), where the invariant features of a signature remain, even when movements are performed with the feet or with the mouth [¶] (Raibert 1977).

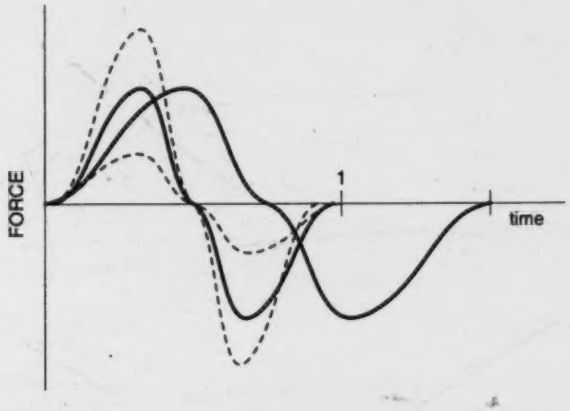

Figure 1.1: The force output of a simple movement. The general observable features of the movement remain constant while the time scale and/or absolute force levels may be varied (Meyer et al. 1982).

1.2.1.4 Response Execution

Traditionally ignored in most cognitive discussions of information processing, the response execution phase involves the 'running' of the prepared program, its update, and correction. At the turn of the last century researchers observed a speed-accuracy trade-off during performance (Woodworth 1899). This result suggested that the human motor system had a limited capacity to transmit information about performance (e.g., Crossman & Goodeve 1983, Fitts 1954, Fitts & Peterson 1964). Paul Fitts (1954) formalized a logarithmic relation between movement time and radial accuracy (Equation 1.3, Fitts' Law). In an aiming task, movement time (MT) was found to vary linearly with the logarithm of the distance to be covered (D) and the width of the target (W).

[¶]Generalized motor programs have been applied extensively to both simple and complex, rapid and long duration movements over the last 20 years.

$$MT = b + log_2(2D/W) \tag{1.3}$$

If compared with the previously discussed Hick's Law, the common unit of bits (binary information) is noticed. Thus the complexity of any movement, both in selection and execution, is dependent upon the amount of information to be processed.

1.2.1.5 The Role of Feedback

Due to the mathematical simplicity of Fitts' Law (Equation 1.3) and its ability to account for situations involving both macro- and micro-movements, many descriptions of the human motor system have been forwarded which justify the relation (e.g., Crossman & Goodeve 1983, Keele 1968, Schmidt, Zelaznik, Hawkins, Frank & Quinn 1979). In general, most accounts follow the basic framework presented by Woodworth (1899). An initial impulse is proposed to carry the limb some distance toward the target. This *ballistic* phase is followed by a corrective phase which allows the amendment of the movement trajectory based on error detection. One point of discussion within all these models however, is the rate at which feedback may be relayed from the receptor (visual or kinesthetic) to the comparator, and back; initial estimates placed this conduction as high as half a second (Woodworth 1899), more recent values fall nearer to 100 ms (Carlton 1981, Carlton & Newell 1993). A shortcoming of these views is that speed-accuracy trade-offs also seem to occur in the absence of feedback and during the ballistic phase. Regardless of the precise mechanism or process resulting in speed-accuracy trade-off's, issues of feedback and movement information dominate research of movement past and present (e.g., Woodworth 1899, Carlton 1981, Elliott, Binsted & Heath 1999, Elliott, Heath, Binsted, Ricker, Roy & Chua 1999, Chua & Elliott 1993).

1.3 Gibsonian Approaches: Branches of Ecological Psychology

The discipline of ecological psychology has the goal of dealing with the general problem of how actions are coordinated with perceptual information (Michaels & Beek 1995). In doing so, ecological psychologists attempt to link (usually mathematically) environmental information with evoked responses, either overt or covert, in a manner that requires little central processing. The idea that information is provided by the environment instead of being sensed and *translated* within cognitive processes is directly at odds with IP theories. There are three directions from which this interaction between information and action have been described: direct perception, kinetic theory, and pattern dynamics.

1.3.1 Direct Perception

Theories of direct perception are based on the concepts of informational flow and object affordances (Gibson 1966, Lee & Reddish 1981, Michaels & Beek 1995). Specifically, the environment of an animal (human) contains invariant energy signatures (i.e., information). This energy may come in the form of texture gradients, rates of optical expansion/contraction, or similar values that pertain to acoustics or haptics. This information energy surrounds the animal and comes in contact with the animal in the form of *flow* (see Kinetic theory). The animal is thought to actively seek temporal and spatial invariants within the informational flow, thus accessing the *affordances* which portray possibilities for action.

The majority of interest in the area of direct perception has focused on time-to-contact (Lee & Reddish 1981, τ) or its time-derivatives (e.g., $\dot{\tau}$, $\ddot{\tau}$)(Bootsma & Mottet 1997, Zaal, Bootsma & van Wieringen 1998, Zaal, Bootsma & van Wieringen 1999). Time-to-contact is a common timing variable for action and perception. If both processes are

stated using this variable, no translation is needed for appropriate coordination/action. The classic example of τ used by Lee and Reddish (1981) is the wing-folding behaviour of gannets (a diving sea bird). The bird was observed to fold its wings prior to contact with the water; the folding action occurred at a constant *time* from the surface, regardless of the bird's velocity. Time to contact is extracted from optical information flow based on the expansion/contraction rate of elements on the retina. When the movement is prepared in similar terms (i.e., time until wing closure) the two events occur with correct timing regardless of absolute velocity. The benefit of such a time-to-contact strategy is that multiple actions may be prepared based on a single optical τ, thus allowing the synchronization of compound tasks.

1.3.2 Kinetic Theory

Kinetic theory is primarily a conceptual method for examining action-perception based upon a circular relation between the application of muscular forces and the generation of information flow. An example of this reciprocity is the connection between force application during walking and τ (time-to-contact) with objects in the environment. As a person is moving through the environment, the amount of impulse[‖] that is applied to the ground directly affects the rate of object expansion on the retina. Alternately, dependent on the rate of expansion of an object, the force to be applied will be dictated. With this method, one identifies mathematical relations to describe the transformations between force fields (physical interaction) and flow fields (information). While beneficial as a thought experiment, extending the theoretical considerations of action-perception, relatively little headway has been made using this approach due to the uncertainty by which the kinematics of movement may be mapped backward to the real world kinetics (i.e., the inverse dynamics problem) (Kugler, Kelso & Turvey 1980, Kugler, Kelso & Turvey 1982, Kugler & Turvey 1987, Michaels & Beek 1995).

[‖]The horizontal component of impulse is the primary influence due to its relation to forward/backward acceleration. However, the vertical component will likewise influence the direction of optical flow by inducing perpendicular acceleration.

1.3.3 Pattern Dynamics (Dynamical Systems Approach)

A third method of research based on Gibson's writings involves the formulation of dynamical equations of motion to capture the coordination patterns within action-perception systems (see Wallace 1996, for an historical discussion and review). These *formal analogies* (Michaels & Beek 1995) to movement are thought to be independent of the structure of the motor (or other) system.

The output of a dynamical model is expressed behaviourally by changes in an order parameter(s) in response to scaling of a control parameter(s). In this framework, a control parameter is a factor which drives the behaviour of a system (e.g., speed, weight, force). A simple example to consider is the effect of pace on gait: when one changes the rate of locomotion one tends to change gait by which one attains that speed (e.g., walk to run). In general control parameters are relatively easy to identify; however their importance and independence from other control parameters must also be examined. Order parameters, also referred to as collective variables, represent the intrinsic *order* of the coordinated system; these factors are very restrained in their definition and must demonstrate specific behaviour.** A collective variable must:

1. display relatively small variations in behaviour in response to large changes in a control variable;

2. exhibit large shifts (phase transitions) in behaviour in response to small changes in a control parameter that satisfy some criterion value (i.e., a critical frequency);

3. show asymmetric behaviour (hysterisis). Specifically, if a phase transition is shown at a critical value of a control parameter while that parameter is increasing, the critical value will differ if the parameter is decreasing; and

**Currently, only one collective variable, relative phase, has demonstrated all the required behaviours of such a variable in human research. Relative phase is the relative timing of positions of movements during a cyclical movement

4. demonstrate critical slowing and critical fluctuations. During phase transitions, behaviour of the collective variable will show a sharp increase in variability (fluctuations) and will become slower at returning to a destabilized state.

While this approach has many proponents (e.g., Byblow, Bysouth-Young, Summers & Carson 1998, Carson 1995, Schmidt & Turvey 1995, Schmidt & Turvey 1994) the seminal work has been completed by Kelso and colleagues (Fuchs & Kelso 1994, Haken, Kelso & Bunz 1985, Kelso 1984, Kelso 1994) and thus will be a focus for discussion in the next chapter.

In this chapter we have introduced some of the general concepts necessary for a the study of human perception-action. The most important feature of any examination of action – as presented by both Bernstein and Gibson – is coordinated interconnectivity between component systems (e.g., perception and action), limiting the requirements on cognitive system. In the next chapter we deal with formal implementations of models for action (and sometimes perception). The selected examples span the motor control literature at a theoretical level, giving a frame of reference within which to approach the development of an alternate framework for action (i.e., one implementing hidden Markov models and their extensions).

Chapter 2

Current Models in Motor Control

Research into human action presents many outlets for applications to medicine and industry. Specifically, a better understanding of the meta-processes guiding the motor control system would aid in the diagnosis and treatment of pathologies (e.g. Parkinson's disease) as well as the design of better working environments (e.g., reduce back strain). One method of developing this knowledge is to attempt the construction of a working system that mimics the behaviour of the human system. The premise of such models is that if a model accurately predicts and generates human-like actions *and* is consistent with known neurophysiology, knowledge regarding the organization of the motor system will be gained.

Presented in this chapter are précis of important classes of models in motor control. These models all attempt to account for human motor behaviour; each however uses differing assumptions (physical, mathematical or biological) as well as having ranging levels of generalizability and unique foci. The discussed models are grouped by their theoretical basis, beginning with the information processing tradition (Meyer, Abrams, Kornblum, Wright & Smith 1988, Rosenbaum, Meulenbroek, Vaughan & Jansen 1999), continuing with engineering-biological hybrid approaches (Bizzi 1980, Plamondon &

Alimi 1997, Uno, Kawato & Suzuki 1989), and ending with self-organizing systems (Bullock, Cisek & Grossberg 1998, Haken et al. 1985).

2.1 The Optimized Sub-Movement Models

As presented in Chapter 1, when people perform goal-directed movements (e.g., pointing) there is an ordered relation between precision and speed. For some time, feedback based models (visual and/or kinesthetic) have been employed to account for speed-accuracy tradeoffs (Fitts 1954) in human movement and in doing so explain Fitts' Law (see Equation 1.3). The deterministic iterative-corrections model (Crossman & Goodeve 1983) and subsequent extensions (Keele 1968, Meyer et al. 1982) consider the act of goal directed aiming (from home position to target in one-dimension) to be composed of a series of discrete submovements, dependent on feedback. In these formulations each submovement is hypothesized to last a *constant* duration (t) and to travel a *constant* proportion of the remaining distance (D) to target centre ($p_d D$). Thus, the distances traveled by the first three submovements would proceed as: $p_d D$, $p_d D(1-p_d)$, and $p_d D(1-p_d)^2$ (Meyer et al. 1988). The movement ends when a submovement terminates inside the target boundary.

In this relatively simplistic manner, the deterministic interative-corrections model can account quite adequately for Fitts' Law. Given $p_d D$ and t, the total time elapsed for the movement would approximate a logarithmic function of distance (D) and target size (W). However, the model is incapable of accounting for many other features of rapid aiming, specifically, spatial variability of submovement endpoints and the frequency of higher-order submovements. The common failing in all these instances is the deterministic nature of the model. The model predicts that for every combination of distance and width there should be a *determined* number of submovements each of a *specified* length and duration, each using feedback to shorten the distance to the target. However, human action is inherently *stochastic*; submovements do not occur in a

stereotyped fashion (Chua & Elliott 1993, Elliott, Binsted & Heath 1999). In addition, aiming, including submovements, is possible in the absence of feedback (e.g., Elliott & Allard 1985, Carlton 1981, Carlton 1992, Chua & Elliott 1993).

Based on the common premise that human reaching is composed of feedback guided submovements, Meyer and colleagues (Meyer et al. 1988) presented the stochastic optimized submovement model to account for many of the shortfalls of its deterministic predecessor. Central to this model is research on neuromotor noise, and specifically, examinations of impulse variability (Meyer et al. 1982, Schmidt, Zelaznik & Frank 1978, Schmidt et al. 1979, Wright & Meyer 1983) which attribute increasing noise in movement to variablity induces by motor-impulse scaling. In addition, the model assumes that an ideal compromise is reached between the lengths and durations of the submovements; the exact form of this criterion is a reflection of neuromotor noise and required precision.

The stochastic optimized submovement model may be parsed into four central assumptions reflecting the inclusion of neuromotor noise into the controlling system:

Number of submovements: All rapid aiming movements are assumed to contain either one or two components submovements, regardless of distance or target size. The primary submovement is programmed to end at the centre of the target, if this is successful, no secondary submovement occurs. If an error occurs, the secondary submovement is thought to correct the error based on feedback (see Figure 2.1).

Distribution shape of submovement endpoints: Based on one of the model's central tenets, neuromotor noise, the distribution of endpoints for both the primary and secondary movements are proposed as having normal (Gaussian) distributions. As a performance correlate of variability, the standard deviation of reaching movements are assumed to increased proportionally to the average velocity of movement.

Movement time minimization: The average velocity of the primary and secondary

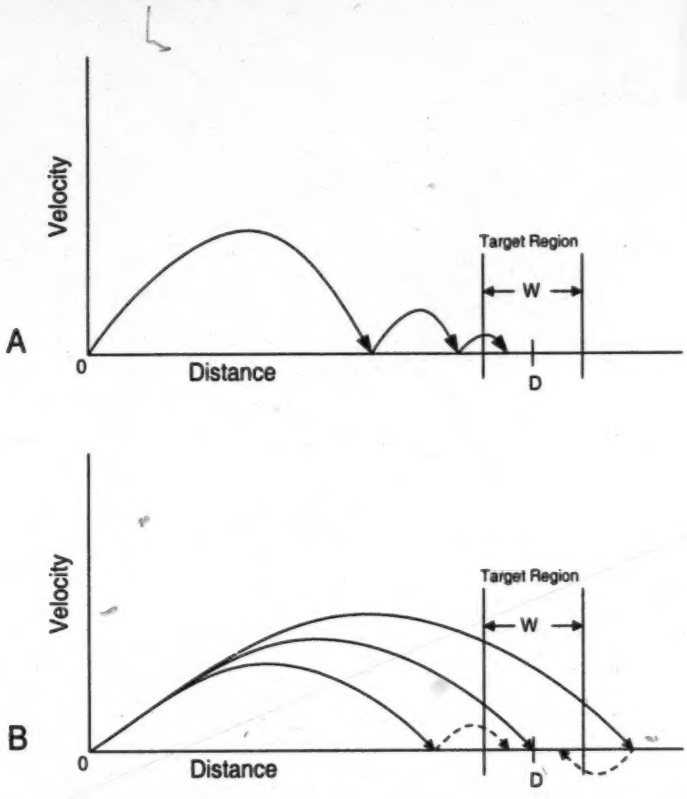

Figure 2.1: A. The Deterministic Iterative Corrections Model (Crossman & Goodeve 1983, Keele 1968). Each submovement is thought to last a constant duration and to travel a constant *proportion* of the remaining distance. B. The Stochastic Optimized Sub-movement Model. Solid lines indicate possible trajectories of primary submovements. The dashed lines indicate feeback guided secondary submovements used to correct errrors in the initial movement. Figure is a recreation based on (Meyer et al. 1988).

submovements is considered to be programmed to minimize average movement time (based on a Gaussian distribution). Optimization occurs over both submovements; as movement velocity increases in the primary submovement, error, and therefore the required distance to be covered by the secondary submovement, increases. Given this error and the goal of accurate aiming, the secondary submovement must proceed slowly enough to guarantee success ($\geq$95%).

Preparatory processing: Information regarding the target distance and size are as-

sumed to be available to the control system during programming of the primary submovement. The secondary submovement is prepared 'on the fly' based on information regarding relative target location and movement kinematics (e.g., peak velocity). Although this transmission of information was originally conceived with feedback as a mechanism, the secondary submovment may equally occur as a result of feedforward processing.*

The stochastic optimized submovement model has found wide support from the motor control research community, thus indirectly validating the model's predictions regarding submovements and neuromotor noise. However, like the previous models in the IP tradition, this model too has limitations. Meyer's model demonstrates remarkable predictive power regarding simple aiming behaviour, however it makes no attempt to attach mechanisms. Moreover, proponents make explicit attempts *not* to indicate anatomical correlates. While many of the assumptions have found support, considerable controversy still surrounds the assumption of a Gaussian distribution for submovement endpoints (Elliott, Binsted & Heath 1999) as well as the notion of planning a movement based upon force/impulse (see Section 2.4 for a review). Finally, both the stochastic model and its deterministic predecessor are very limited to goal-directed aiming and prove difficult to 'scale-up' to account for more complex actions.

2.2 The Knowledge Model

Continuing in the IP tradition, David Rosenbaum and colleagues (Rosenbaum, Engelbrecht, Brushe & Loukopoulos 1993, Rosenbaum et al. 1999) have recently begun examining human grasping using a "biologically and physically plausible" (Rosenbaum et al. 1993, pp 238) approach. Unlike Meyer (Meyer et al. 1988), no attempt was made to explicitly account for dynamic features of movement. Instead, the reaching control

*Feedforward is defined here as encompassing efference copy and corrollary discharge generated during the programming/preparation of the primary submovement

task was regarded as a problem of selecting appropriate joint postures with end-point position emerging as a result. The *Knowledge Model* is thus cast as a solution to a "multiple-constraint satisfaction problem" (Rosenbaum et al. 1999) for solving the degrees of freedom problem in a redundant human limb.

Overall, the Knowledge model is a rule-based approach to human reaching. Graphically, the Knowledge model is quite simplistic; a seated human is represented as a rendered stick-figure with three degrees of freedom (hip, shoulder, elbow) able to move only in the sagittal plane. The model is tested in a *Turing*-like manner,[†] requiring the model to reach to a region of space within its 'grasp'.

2.2.1 Storage

The base element of the Knowledge model is posture, with actions being the transitions from the present posture to stored postures. A posture (joint angle of multiple joints) is encoded by registering the Cartesian location of the hand resulting from attainment of the required joint positions. By reducing the model to the expression of posture states, Rosenbaum enables a reduction in the required storage *space* to the dimensionality of the component joints. Further, a *mass-spring* (Bizzi, Ivalidi & Giszter 1991) mechanism for movement is employed, with the transition between postures occurring in a manner corresponding the physical properties of the system (i.e., spring-like properties of muscle).

2.2.2 Planning

The process of planning in the Knowledge model is entirely driven by the evaluation and minimization of the *total* cost for attaining each sub-posture in a given reaching

[†]A Turing test is a test of intelligent machine, determining the degree to which the artificial system mimics human behaviour. Specifically, a person is set with the task of determining whether the responses being generated (in this case a movement) is the output of a machine or a human.

movement. Total cost is considered the sum of the *spatial error* cost and the travel cost. The spatial cost is simply equated to the squared Euclidean distance between the Cartesian location of the target and proposed location of the hand if the posture was selected. Travel cost, V_p, for each posture is defined as:

$$V_p = \sum_{j=1}^{n} \xi_j \left(\frac{\alpha_j}{t} \right) \tag{2.1}$$

where α_j is the distance in joint centered coordinates, for the j^{th} joint of the n required to define the posture.[‡] Movement time, common for all joints, is expressed as t, with ξ_j denoting the *cost factor* associated with the movement.

The cost factor, ξ_j, is a representation of the friction, intertia, damping, and stiffness associated with the ensemble segments. In general, this factor scales with moment of inertia for each joint and is assumed to remain invariant for a segment's contribution to all postures. However, the general cost definition of this variable also enables it to be increased for situations such as injury or muscle stiffness (Rosenbaum et al. 1993, Rosenbaum et al. 1999).

In the original formulation of the Knowledge model, goal posture was determined by combining the total cost, V_p, of all prospective postures into a weighted sum (Gaussian average) using total cost as a scaling factor (Rosenbaum et al. 1993). By assessing potential postures in this fashion, previously unattempted postures may be selected by the motor system. [§]

A more recent version of the model takes a somewhat different approach to posture selection (see Figure 2.2); goal postures are selected in a two-stage process based upon the last N postures attempted (limited memory). A search of these *memory* postures is completed based upon a *constraint hierarchy* (Rosenbaum et al. 1999), containing both

[‡]Joints in the Knowledge model were considered to only consist of one degree of freedom therefore dfs and joints may be considered synonymous.

[§]The Gaussian function is shaped such that the mean is zero and the standard deviation is proportional to the smallest total cost.

implicit and explicit requirements of the task (i.e., more weight may be placed upon velocity requirements than accuracy) to be used as gain factors for potential postures. Once a likely posture solution is selected, a local search is completed from the postures *neighbouring* the selected one, to see if requirements may be better satisfied (i.e., results in a lower total cost).[¶] The search is continued in an expanding fashion (number of degrees away from memory position) until a pre-specified temporal deadline is reached. This deadline may also be updated based upon the whether a criterion satisfying posture was found prior to the deadline (reduce time) or had not yet been found (increase time).

2.2.3 Execution

Execution of movements proceeds according to an assumption of a common movement time for all joints. While this time-stimulation may seem overly contrived, such a temporal scenario is required for linear motion to be achieved by the end effector. Moreover, several examples have been identified for discrete actions where this feature of termination timing exists (e.g., Carnahan & Marteniuk 1991, Carnahan & Marteniuk 1994, for eye-hand-head coordination in reaching). Thus, within each posture transition, movement dynamics may defined such that:

$$\omega_j(t) = \eta_j \omega_j(t-1) + (1-\eta_j)\omega_j^*(t) \tag{2.2}$$

where $\omega_j(t)$ is the angular velocity of the j^{th} component at time t, and η_j is proportional to the inertia of that element.[*] The final term, $\omega_j^*(t)$, represents the *intended* angular velocity of the joint component and is directly proportional to the total angular distance to be covered by that component's motion. An important consideration in this representation is its formalization of feedback processes (terms to the left of the

[¶]These neighbouring possibilities include postures with slight variations in individual joint angles from the stored positions.

[*]The term η_j is a dimensionless *index* of inertia ($0 \leq \eta_j \leq 1$).

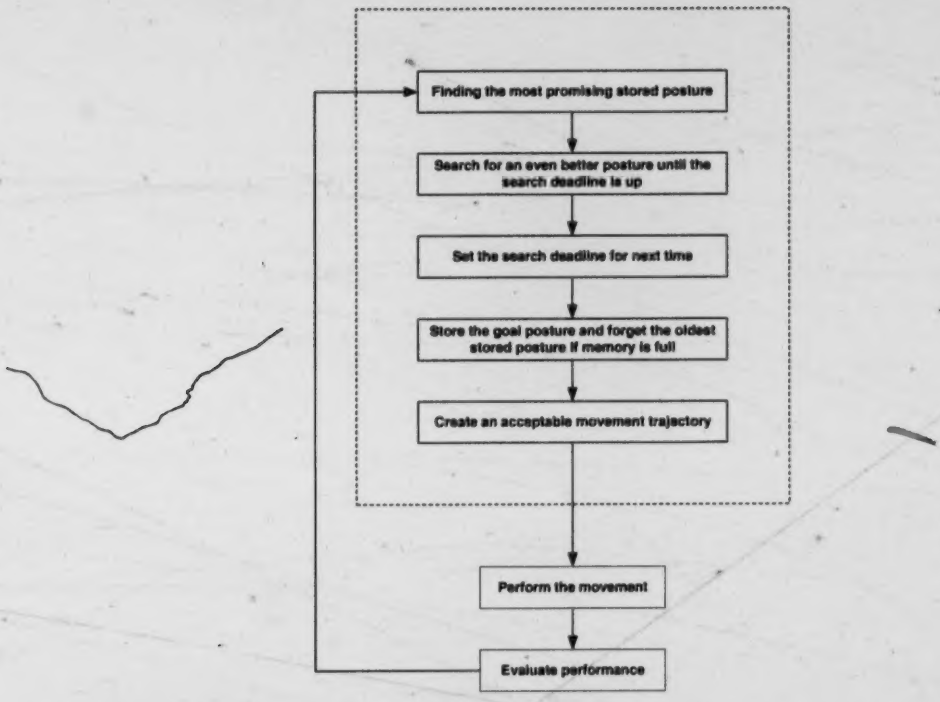
Finding the most promising stored posture
Search for an even better posture until the search deadline is up
Set the search deadline for next time
Store the goal posture and forget the oldest stored posture if memory is full
Create an acceptable movement trajectory
Perform the movement
Evaluate performance

Figure 2.2: A schematic representation of the Knowledge Model (Rosenbaum et al. 1999). Each box represents a time-occupying event within the planning process. Once the movement to the selected posture is completed, a feedback-loop updates the process with a new hand position until the hand and target position coincide.

Knowledge figure is recreated from (Rosenbaum et al. 1999).

addition sign).

The Knowledge model demonstrates some features important for any model of human action. In particular, the model easily accounts for movement in the absence of prior experience; a search is completed with a start posture from memory (of *similar* tasks). Advantages are also gained by prior knowledge (i.e, if a memory posture satisfies the cost requirements). The model accounts for intangible costs associated with normal human movement. Quite simply, if a joint or muscle is injured, all cost associated with motion requiring that item increases. If a component is completely immobilized, movement may continue by assigning a high cost to postures involving that segment.

This approach is also perhaps the most behaviourally complete model of reaching, accounting for wide variations in tasks and system geometry/features. Despite demonstrating qualities that satisfy a limited Turing-like criterion, the model fails on several fronts. The Knowledge framework is based upon *cost* optimization, specifically, a sum of quantitative (energy) costs and a subjective factor to account for pain, obstacles, and other similar features (see Section 2.3). The second, subjective cost, is both the strength and weakness of the model. This factor enables the model to account for a wide variety of situations using a relatively simple selection/assessment criterion as well as allowing the description and prediction of pathological movements based on subjective and clinical observations. However, even Rosenbaum and colleagues acknowledge lack of an explicit definition of the cost, i.e., ascribing the cost to a user defined term(Rosenbaum et al. 1993). Added criticism should also be made due to the vagueness of even the quantitative measures. These factors, while more objective, also lack an explicit definition. The inertial cost of moving various limbs is determined only by an ordinal estimation of joint moments. †

2.3 Minimum Torque-Change Model

Moving toward a purely computational approach to movement production, Uno and colleagues (Uno et al. 1989) proposed a model based on minimization of energy related parameters as a solution for the complex task of trajectory planning during rapid two-dimensional aiming. Unlike the Knowledge model which attempted to pose the movement problem from an information processing standpoint, the Minimum Torque-Change approach is primarily an 'engineered' solution to the indeterminacy problem (with biological constraints). As such, the necessary cognitive processes for this system

† Cost is assigned based upon rotating weight. Therefore a hip movement would be assigned higher cost than shoulder, which in turn would be higher that elbow and wrist. However, the relative sizes of these values is indeterminable.

are: (1) select a appropriate trajectory[‡], in extrinsic coordinates, from the infinite possibilities (see Chapter 1), (2) transform the coordinates of this trajectory into an egocentric reference frame, and (3) select the appropriate motor commands to generate the needed torques about the component joints. The minimum torque-change (MTC) model approaches the reaching problem more completely than some previous energy minimization approaches in that multi-joint movements are considered (Flash & Hogan 1985).

Although demonstrably superior in performance, the MTC model is based on the minimum-jerk (MJ) model proposed by Flash and Hogan (1985). This MJ model was formulated based on the observation that a unique (ideal) trajectory is defined by the minimization of an objective function (Equation 2.3). The function is defined by the square of the jerk (rate of change of acceleration) of the end effector, integrated over the entire movement, during a movement of duration, t_f,

$$C_J = \frac{1}{2}\int_0^{t_f}\left\{\left(\frac{d^3x}{dt^3}\right)^2 + \left(\frac{d^3y}{dt^3}\right)^2\right\}dt \tag{2.3}$$

where (x, y) are the Cartesian coordinates of the hand. Based on the minimization of Equation 2.3, there is an expectation for a straight line path between the origin and the end location, and for a bell-shaped velocity profile (Uno et al. 1989). Correspondingly, Flash and Hogan (1985) demonstrated that this model successfully predicted several qualitative and quantitative features of reaching movements performed in front of the body. However, as Uno et al.(1989) point out, given that the movement is defined only in terms of start/end position and kinematics, there is no place for quantities such as *mass* and *torque* in the determination of trajectory.

The MTC model is based upon a very similar objective function to that proposed in the MJ formulation. In the MTC model however, specifications of mass, selected muscle,

[‡]Trajectory is defined as both the path and the dynamics of the movement

and desired torque are included (Equation 2.4).

$$C_T = \frac{1}{2}\int_0^{t_f} \sum_{i=1}^{n} \left(\frac{dz_i}{dt}\right)^2 dt \tag{2.4}$$

The motor command (torque, z_i) is sent to the i^{th} of n muscles such that the criterion (C_T) is the sum of the squared rated of change of torque integrated over the movement (time, t_f). While some properties of the function are similar to that in the MJ model, the prediciton of a trajectory requires consideration of the structure and dynamics of both the environment and the musculoskeletal system. Due to the overwhelming complexity of these systems, simulations are performed with a system containing a two-joint robotic limb constrained to two dimensions. The limbs are also controlled by only two muscles presumed to be acting at known lengths, rotary inertia, etc. Even given the severe limitations placed upon the system, no analytically obtainable solutions were forthcoming.[§]

As mentioned previously, the MTC model provides superior results to the MJ model (Uno et al. 1989). When short movements were simulated as occurring approximately in front of the body, results remained similar to the MJ model, qualitatively and quantitatively. However, when movements were simulated in situations which either departed from the mid-line of the 'subject' or were large in amplitude however, performance was different. Under the MJ model, trajectories remained relatively straight. Conversely, the MTC model successfully predicted a convex curve in the motion commensurate with the region of workspace and movement scale. Benefits were also shown for movements which required a via-point to be attained (Uno et al. 1989).[¶] The MTC model predicted adaptations in movement speed profiles dependent not only on the movement geometry but also on the segment postures and external force.

Although the MTC and MJ models demonstrate substantial predictive potential for

[§]The researchers resolved the minimization through an iterative learning procedure (Uno et al. 1989).
[¶]A via-point is merely a required intermediate point which requires trajectory adaptation to attain.

simple movement trajectory formation, they have common failings. Both models require the motor system to predetermine the entire path and dynamics of the movement prior to execution. While this assumption allows for energy minimization and thus ideal path trajectory selection, it ignores a substantial literature based on human trials that demonstrates continuous, or on-line, control of reaching (for example Chua & Elliott 1993, Elliott, Binsted & Heath 1999). A second shortfall is the reliance upon over-simplified motor systems. Due to the requirements of the MTC and MJ models to quantify then derive parameters, the computational load required grows exponentially with system complexity. In addition to the limitations that simplification places on simulation validity, there are significant drawbacks generated by the high hypothesized loads upon the human cognitive system (c.f., Gibson 1966, Gibson 1979).

2.4 Equilibrium-Point Models

At a similar level of description to the Minimum Torque Change model, a series of models exist which describe human movement based on the equilibrium (static or dynamic) emerging from the interaction between muscles. Known as the mass-spring models (Asatryan & Feldman 1965, Bizzi, Accornero, Chapple & Hogan 1982, Bizzi et al. 1991, Feldman 1966a, Feldman 1966b), this general class of approach is based upon the anatomical properties of human neuro-muscular system (Figure 2.3) and seek to account for control based on the expression of these features. Specifically, muscle is proposed to contain a contractile element (T_0) as well as damping (dash-pot, b) and elastic elements (k_1 and k_2)(McMahon 1984). In Figure 2.3, the elastic elements are connected in series to each other and in parallel to both the damping and contractile elements. This configuration results in the muscle behaving in a fashion similar to a damped (non-linear) spring whose output force is dependent upon length. This force-length relation is precisely the property that mass-spring models exploit in their explanations of the human movement control process. These models also integrate the

known physiology of the nervous system, using reflex mechanisms specifically associated with tissue length, tension, and velocity.

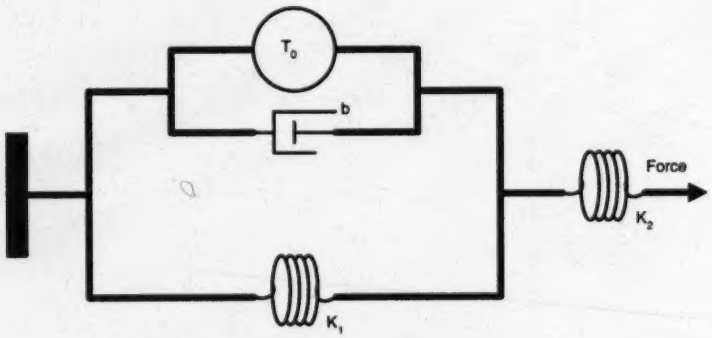

Figure 2.3: A muscle represented as a composite of spring, contractile, and damping elements. These feature connected in series and parallel replicate a large proportion of muscle behaviour as well as coinciding well with the physical stuctures involved.
(From Latash 1993)

2.4.1 The γ- and α-Models

In 1953, P.A. Merton (Merton 1953, Eldred, Granit & Merton 1953) provided the first of the mass-spring models by including the load compensation properties of the tonic stretch reflex into a comprehensive model of control. This model, the γ-model, has as its central tenet that load applied to a limb stimulates a stretch reflex thus modulating α-motoneuron pool activity and generating a firing of muscle agonists. The likelihood of this reflex occurring is set within the γ system by central specification of muscle spindle sensitivity ranges and thresholds (Latash 1993, Kandel, Schwartz & Jessell 1991). This model has high appeal due to its simplicity; a central motor command specifies a new *position* (muscle length) within the γ system, thus stimulating a reflex response and generating closed-loop agonist activity to attain a position.

In the tradition of Merton, E. Bizzi and colleagues (Bizzi 1980, Bizzi et al. 1982, Bizzi et al. 1991) generated the α-model. While borrowing the *servo* architecture from the γ

model, the α-model placed the role of response modulation at the α-motoneuron pool level. Specifically, descending inputs fix the transient activation of the pool for a given muscle, thus determining the amplitude of response for a given load (as determined by muscle spindle firing). Even though the α-model has received considerable criticism for inconsistencies with regard to experimental evidence (Feldman 1986, Berkinblit, Feldman & Fukson 1986), it still remains a point of discussion in the literature due to its compelling prediction of electromyogram (EMG) firing patterns in muscles and their direct relation to α-motoneuron pool activity.

2.4.2 The λ-Model

The λ model proposed by A.G. Feldman (Feldman 1966a, Feldman 1966b, Latash 1993) has many common features to other mass-spring models, at a fundamental assumptions level. The name of this model is derived from length threshold (at a muscle spindle level) necessary to generate recruitment of α-motoneurons during tissue stretch, λ. A central motor command is proposed to define λ, thus choosing a force-length curve (Figure 2.4a). Actual muscle length however is also defined by external load. In Figure 2.4b, three forms of external load are represented (Latash 1993): (1) isometric, (2) isotonic, and (3) elastic. The point at which load intersects the force-length curve will determine the equilibrium point of the programmed movement. In this manner, the spring properties are paramount to the λ-model (see Figure 2.3), allowing the limb to shift position simply by establishing of new "equilibrium point"(EP).

By utilizing the spring-like properties of the muscles, the central nervous system allows the specification of a new location for a limb in a relatively simple manner. Consider a joint defined by only one degree of freedom, and, controlled by a single pair of muscles (agonist, antagonist). In order to specify a movement, the task is to specify the net torque about the joint by constraining λ of each muscle given load (L). The

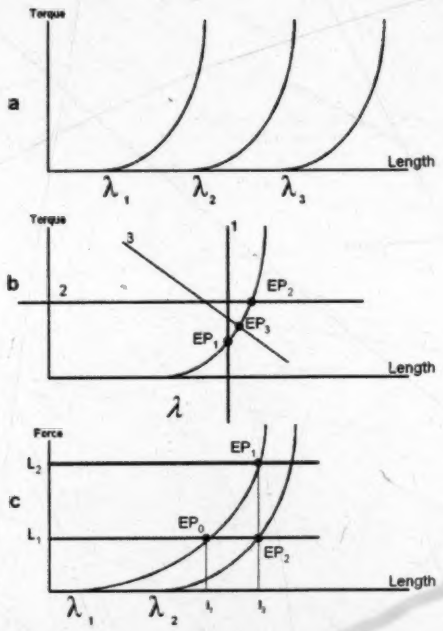

Figure 2.4: (a). Examples of how simply by choosing a threshold λ a descending command also specifies a force-length curve. In order to determine an equilibrium point (EP) a load must be present. (b). Examples of possible load parameters (1) isometric, (2) isotonic, and (3) elastic. Each load results in a unique EP for a given λ. (c) An example where a a new EP (l_2) is attained by a shift in load or by the selection of a new force-length curve (λ) (Latash 1993)

command will select a new EP for each muscle thus determining its activity state.† The muscles will continue to be active, according to the relative EP, until the new location (equilibrium) is attained.

One strength of equilibrium-point models is their attention to mechanisms for the micro-control of limb segments. Control is accomplished in a fashion that is entirely joint-centered, and is therefore invariant to alterations in absolute space. As such, the γ, α, and λ models account well for the considerable number of findings that demonstrate *equifinality* (for example, Bizzi, Polit & Morasso 1976, Latash & Gottlieb 1990, Polit & Bizzi 1978), a scenario where a limb is perturbed unexpectedly by a transient load

†If a muscle's length exceeds that specified by the EP then the muscle is *active*.

yet still reaches its equilibrium-point. However, although the mass-spring hypotheses utilize many mechanisms based on the known physiology, they encounter difficulties when fully complex actions are considered. Specifically, no comprehensive spring-based formulation has been produced that easily predicts the movement and control processes during even simple multi-joint movement. The complications arise due to the relation between joint torque and 'length'. A servo-controller based on these combined factors has difficulty accounting for multi-joint muscles, as length will be modified due to neighbouring joint rotation. This shortfall is important, as all 'real' movements require control of torques about multiple joints, even if the joints must remain immobile. An added problem associated with an equilibrium point process of control is the extensive behavioural research supporting a force-specification approach (e.g., Abrams, Meyer & Kornblum 1989, Binsted, Cullen & Elliott 1998, Elliott, Heath, Binsted, Ricker, Roy & Chua 1999, Schmidt et al. 1979, Sherwood 1986). Force or impulse based models easily account for discrete/continuous corrections (Chua & Elliott 1993) as well as adaptive behaviours such as target undershooting (Elliott, Binsted & Heath 1999). Conversely, an equilibrium-point model generates the expectation that movement endpoints should be equally distributed about the target. There is also the false prediction that if a muscle behaves as a spring, the limbs would be expected to show endpoint oscillations. In addition, force specification may be easily matched with the characteristic speed variability curves of movement and the expected triphasic – agonist-antagonist-agonist – EMG pattern for discrete action.

2.5 A Kinematic Theory of Rapid Human Movement

Returning again to the discussion of the speed-accuracy trade-off, R. Plamondon and colleagues (Plamondon 1995a, Plamondon 1995b, Plamondon & Privitera 1995, Plamondon & Alimi 1997, Plamondon 1998) developed a kinematic theory based on the impulse response of neuromuscular systems. While this approach retains much of the

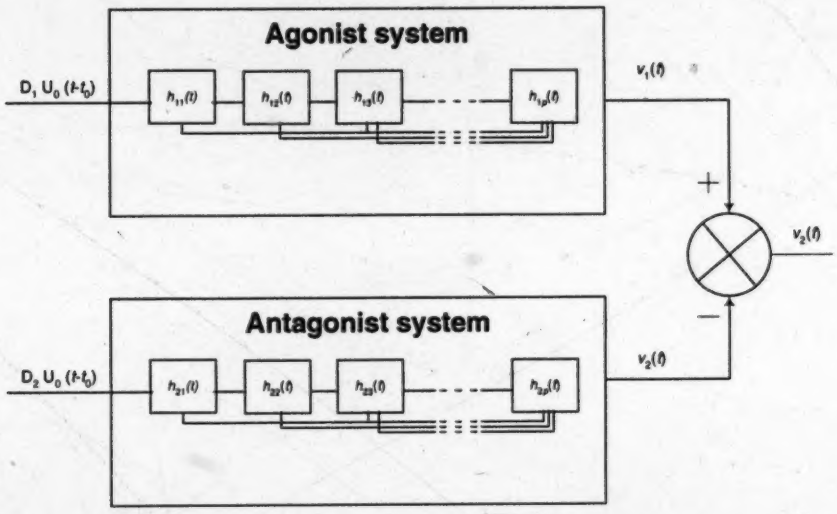

Figure 2.5: The Kinematic Theory. An impulse driven representation of a two muscle synergy. The command, and its amplitude, are specified to each muscle with the output being the result of cumulative impulse response of the contributing system. Movement velocity occurs if the difference between the agonist and antagonist outputs reaches threshold. Figure is a recreation based on (Plamondon & Alimi 1997).

same 'engineered' feel as the Minimum Torque approach, it makes a considered (and successful) attempt to integrate many of the muscle properties discussed in the Λ model. Unlike other approaches however, the central point of Plamondon's proposal is that speed-accuracy trade-offs are inherent constraints arising from the impulse response of the neuro-muscular system. This system is assumed to be composed of a *synergy* of two parallel systems (agonist and antagonist), each a composite of many subsystems (see Figure 2.5).

Several initial assumptions are made. First, due to the relative invariance and theoretical importance of kinematic profiles in extant studies of rapid aiming, showing an asymmetric bell-shape (for example, Chua & Elliott 1993, Elliott, Binsted & Heath 1999, Uno et al. 1989), outputs of the system were modeled based on limb velocity profiles. Second, due to the intent of describing well-learned action (simple aiming), subsystems were considered to act linearly and to operate in a primarily feedforward fashion. As

such, according to kinematic theory, both agonist and antagonist systems produce a velocity output ($v_1(t)$ and $v_2(t)$ respectively) based on an impulse program ($U_0(t-t_0)$) with amplitudes D_1 and D_2 (Equation 2.5).

$$\begin{aligned} v(t) &= v_1(t) - v_2(t) \\ &= D_1 H_1(t-t_0) - D_2 H_2(t-t_0) \\ &= D_1 \Lambda(t; t_0, \mu_1, \sigma_1^2) - D_2 \Lambda(t; t_0, \mu_2, \sigma_2^2) \end{aligned} \tag{2.5}$$

where

$$\begin{aligned} \Lambda(t; t_0, \mu_i, \sigma_i^2) &= H(t-t_0) \\ &= \frac{1}{\sigma_i \sqrt{2\pi}(t-t_0)} \exp - \left\{ [\ln(t-t_0) - \mu_i]^2 \cdot \frac{1}{2\sigma_i^2} \right\} \end{aligned} \tag{2.6}$$

The output velocities may also be viewed as the difference between two log-normal functions ($D\Lambda$, Equation 2.6). In these functions μ refers to the total log-time delay ($h_{i1\ldots ip}(t)$) and σ to the log-response time across all elements of the system , relative to t_0. In summary, the execution of a rapid aiming movement may be viewed as resulting from the synergetic activation of agonist and antagonist systems (at $t = t_0$). The systems respond to their respective D_i amplitude commands with an a log-normal response (Λ, Equations 2.5, 2.6) based on the system parameters σ_i and μ_i.

Referred to as the $\Delta\Lambda$ Law due to its reliance on the difference between two amplitude weighted log functions, this relation (Plamondon 1995a, Plamondon 1995b, Plamondon & Alimi 1997) has been used to replicate a variety of experimental findings. The theory reproduces single, double, and triple peaked velocity profiles based only on a single synergy, using no feedback. This result is contrary to a large body of literature which attributes such velocity features as representative of feedback processes (Carlton 1981, Chua & Elliott 1993, Chua & Elliott 1997, Elliott, Binsted & Heath 1999). In addition to this initial departure, the kinematic model has been criticized for its lack of a viable theoretical framework (Bootsma & Mottet 1997), specifically with regard to

the feature of 'sensory maps' for movement preparation (Carlton & Liu 1997). A final shortfall of the model is its apparent inability to account for motor variability (Jong & Galen 1997); while there are no explicit 'noise' terms, the model does draw upon Gaussian distributions of σ_i and μ_i to generate the impulse response.

Although this approach has several theoretical weaknesses, it explicitly addresses many of the key features necessary for a comprehensive model. Specifically, the model is inherently stochastic and dynamical – features which will be at the forefront of later discussions of the Bayesian modeling approach to action. Also, the control predictions made by the model are based on emergent properties and relationships within the physical system, not relying on an all-knowing homunculus, a property at the core of the next models in this discussion.

2.6 The *VITE* Model

From the previous discussion it is clear that there is a tendency for current models to either attend to the cognitive and psychological aspects of control *or* the physical/anatomical aspects via computational loading – but never both. However, the vector integration to endpoint (*VITE*) model of via-point movements presented by Bullock and colleagues (Bullock & Grossberg 1988, Bullock & Grossberg 88, Bullock et al. 1998, Bullock, Bongers, Lankhorst & Beek 1999) is a departure from previous models. That is, the model is one of a collage of models presented by these authors to describe a wide range of behaviours, accounting for a large variety of kinematic and psychophysical phenomena while remaining consistent with known neurobiological events and structures. The *VITE* model makes explicit predictions of how learning and performance rely upon feedback while remaining relatively free of memory constraints (Bullock et al. 1999). The most recent elaboration has been explicitly applied to the problems associated with voluntary neural control of 2-dimensional arm trajectories.

Owing to a Gibsonian (kinetic theory) standpoint, the *VITE* model for reaching stresses the idea that information detection and motor performance are mutually dependent processes (see Figure 2.6). In general, the *VITE* circuit is composed of the processing required for the preparation, execution, and correction of point–to–point movements. In order to calculate a movement vector to guide the hand to a desired position, both a target position vector (TPV) and a present position vector (PPV, i.e., current position of the hand) are generated in body-centered space. A difference vector (DV) is then calculated with respect to these two positions, specifying both magnitude and direction of movement required to contact the target. The DV is gated by a scalable GO parameter (i.e., "start"); the GO signal is set to zero during movement preparation, thus preventing execution, and scaled to generate voluntary changes in movement kinematics independent of DV. In addition to permitting motor priming operations, the GO signal provides a convenient mechanism for movement equifinality (for example, Bizzi et al. 1976, Latash & Gottlieb 1990, Polit & Bizzi 1978).

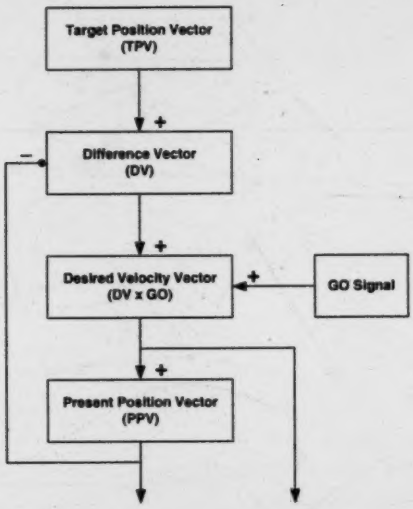

Figure 2.6: A demonstration of the *VITE* circuit. Initially a difference vector is calculated between the current position and the target. This vector is scaled by the GO signal to determine the dynamics of movement prior to a negative–feedback updating the difference vector. (Bullock & Grossberg 1988, Bullock et al. 1999) Figure is a re-creation based on (Bullock et al. 1999)

The *VITE* representation of movement also makes strong claims regarding cerebral organization and cell types used in via-point movements (Bullock et al. 1998).* In support of this structural view, cells behaving in a manner consistent with DV computation have been identified with the primate cortex (precentral areas 4 and 6, postcentral area 5)(Bullock et al. 1998). Evidence from the visual system provides support for an internal GO mechanism (Grossberg, Roberts, Aguilar & Bullock 1997) residing in the substantia nigra and globus pallidus, gating the output of eye movements from the superior colliculus (e.g., Georgopolous, Kalaska, Caminiti & Massey 1982, Georgopolous, Schwartz & Kettner 1986, Schwartz, Kettner & Georgopolous 1988). This evidence for a GO controller in vision is bolstered by evidence from patients with Parkinson's, indicating a gating role for the basal ganglia in gross-motor generation (Jankovic & Tolosa 1993).

The *VITE* model provides a concise explanation of many behaviours demonstrated both at a kinematic (Bullock et al. 1999) and cellular level. The model, however, makes no attempt to deal with the complexities involved with generating the called for vectors. However, Bullock et al. (1999) do accent the benefit of continuous updating of vectors within *VITE*. While this initially seems a strength of the model, serious concerns may be raised as to the biological feasability of feedback at the rate required to maintain this circuit. Moreover, as is demonstrated by the λ and Minimum Torque models, representing muscular force generation even within a two-muscle system is an odious task, even when done off-line.

2.7 The Haken-Kelso-Bunz Model

The final model, HKB, again features aspects of self organization, not dissimilar to the neural network approach of Bullock and colleagues (1999). However, the features

*Again, via-point actions are those that require some intermediary target attainment prior to reaching the final goal

of this system are somewhat more abstract and focus at a much more macroscopic scale. For example, even casual observations result in the realization that even simple movements require precise temporal and spatial coordination between limb segments. As was emphasized by previous discussions of the degrees of freedom problem (Bernstein 1967), the theoretical computational load to reduce this problem are immense. So, consider now the complex act of walking, and extend the problem to animals with more than two legs (maybe six-legged ants). It is not reasonable to attribute the complex coordination required for multi-pedal motion to the cognitive prowess of an insect, instead let us consider a dynamical pattern perspective and look for features suggesting self-organization.

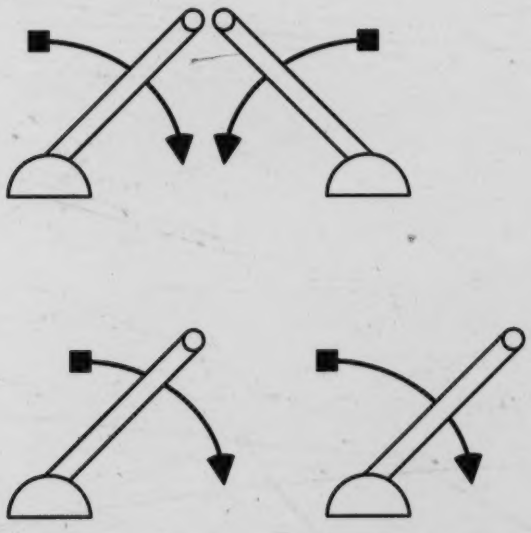

Figure 2.7: Movements of segments may be represented by their relative timing/position. In the upper diagram, the segments are moving in an in-phase fashion, normally associated with concurrent firing of homologous musculature. In the lower diagram, anti-phase is represented.

Borrowing from non-linear dynamics Haken, Kelso, and Bunz (1985) (HKB) proposed a model to describe oscillatory movement coordination (e.g., walking) in terms of a self-organizing system. The experimentation that led to the HKB-model involved the coordination of rhythmic finger movements. Participants were instructed to move their

index fingers (or hands) in either an in-phase or anti-phase manner (see Figure 2.7) [†]. For subjects who prepared their movements in an anti-phase fashion, as the frequency of the movement increased the stability decreased, and an abrupt shift to the in-phase pattern was shown (phase-transition, see Figure 2.8). However, if the movement was prepared in-phase, the coordination remained relatively stable and showed no transition. Haken et al. (1985) proposed a model to describe this behaviour based upon concepts in synergetics (Haken 1977, Haken 1983) and nonlinear oscillator theory.

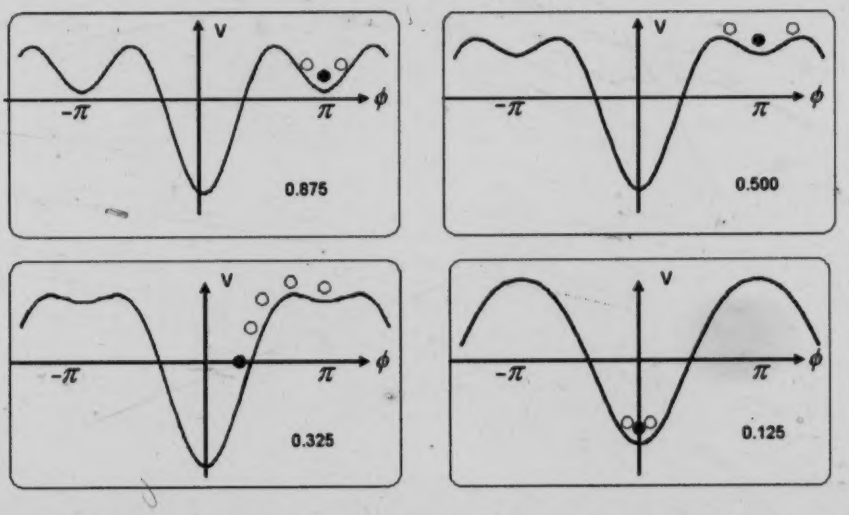

Figure 2.8: All panels represent a potential *landscape* of the system based on the ratio of b and a parameters (value in each panel). Current behaviour is represented by the dark circle and recent behaviour by the open (distribution represents noise). In the upper left panel, relative phase is stable in the anti-phase relation ($\phi = \pi$). As the a/b ratio decreases (top right and bottom left) behaviour becomes unstable and undergoes a phase transition. The bottom left panel has behaviour again stable, at in-phase ($\phi = 0$) (Beek et al. 1995, Haken et al. 1985). Figure is a re-creation of (Beek et al. 1995).

As discussed in Chapter 2, the relative phasing (ϕ) between segments has been identified by dynamical systems theorists as an order parameter for coordinated movements. Thus, the goal of the HKB-model is to accurately predict/describe the behaviour of ϕ across variations in the control parameter (frequency of oscillation). Based upon

[†]In-phase movements involve the concurrent activation of homologous muscle groups; anti-phase, the simultaneous action of non-homologous groups. For example, the alternating action of walking is an ecological demonstration of anti-phase. Similarly, two-handed reaching-grasping is an in-phase movement.

the relatively simple behaviour displayed in the "finger-wiggle" experiment, the system was modeled as containing fixed point attractors at $\phi = 0$ and $\phi = \pi$. In addition, the formulation must be 2π-periodic and symmetric when transformed from $\phi \rightarrow -\phi$ to account for left/right symmetry between coordinated segments and reversals in the leading segment (Beek et al. 1995). Given these constraints and a premise of simplicity, the system was modeled as the superposition of two cosine functions (Equation 2.7).

$$V(\phi) = -a\cos(\phi) - b\cos(2\phi) \tag{2.7}$$

Equation 2.7 adequately describes the behaviour in the finger-wiggle experiment, as well as a considerable number of other experiments and tasks quite effectively(e.g., Carson 1995, Swinnen, Beirinckx, Meugens & Walter 1991, Verschueren, Swinnen, Cordo & Dounskaia 1999). The cosine functions generate minima at $\phi = 0$ and $\phi = \pi$, while the depth of these minima (degree of attraction) is scaled to the ratio of b/a (see Figure 2.8). As the ratio ranges from 1.0 to 0.0 the behaviour of system moves from its initial bi-stable state to a mono-stable one; this ratio may be extrapolated to the frequency parameter originally manipulated by Haken et al. (1985). According this model, once the system has made a transition to the $\phi = 0$ attractor, it will not return to anti-phase (see Figure 2.8).

While the HKB-model satisfies many of the observations of the original experiment (Haken et al. 1985), it fails to represent two key observations in human data. The model is unable to represent the variations in the critical frequency that generated phase transitions. The model is also unable to account for asymmetries in the frequency of transitions initiated by the left and right hands (Byblow et al. 1998, Schmidt & Turvey 1995). As such, two terms have been added (Fuchs & Kelso 1994, Schmidt & Turvey 1995).

$$V(\phi) = \delta - a\cos\phi - b\cos(2\phi) + \sqrt{Q}\zeta_t \tag{2.8}$$

The first, δ, is a detuning factor to induce asymmetries associated with handedness or

similar divergences from the original $\phi \rightarrow -\phi$ assumption (Equation 2.8). The other, $\sqrt{Q}\zeta_t$, is a compound term where ζ is a Gaussian white noise process that functions as a stochastic force of strength Q (Fuchs & Kelso 1994, Schmidt & Turvey 1995). Upon the addition of these terms, the HKB-model has become the closest approximation to a "phenomenological law of interlimb coordination" (Schmidt & Turvey 1994), successfully describing the behaviour of complex biological systems without making any assumptions regarding the mechanistic architecture.

In a recent extension to the HKB-model (Zaal et al. 1998, Zaal et al. 1999), the coordination of prehension movements was proposed with the same dynamical systems framework (for discrete aiming, see also Schoner 1990). Discrete trajectory formation was modeled as 'emerging' during a phase transition from a point attractor placed at the start location to one at the end-goal. By modeling reaching in this manner, the complexities of trajectory selection are accounted for by the 'self-organizing' nature of the HKB framework. Also, in an ecological psychology tradition, the grasp phase of prehension was proposed to be *triggered* in response to time to contact (τ, see Chapter 1). This τ-theshold was based on maximum hand-closing rate during reaching. Although much of this extension may seem *ad-hoc*, the basic ideas of the model (i.e., having the trajectory be defined and controlled by the properties of the system) parallel the mass-spring models.

The fundamental goal of dynamical systems researchers is to develop analogies to movement. While the HKB-model may point to some control heuristics that guide the controlling system, models developed under the umbrella of dynamical pattern theory rely heavily on abstraction. Prior to any behavior being examined with this approach, the task must be *converted* to a form that can be represented by relative phase. Similarly, although there have been recent attempts to use the dynamical systems approach to studies of discrete movement (Schoner 1990, Zaal et al. 1998, Zaal et al. 1999) they remain founded on a literature of cyclical actions and associated assumptions. Finally, although the dynamical approach has been applied to a wide variety of tasks like walking

(Diedrich & Warren 1998, Warren, Kay & Yilmaz 1996), reaching (Zaal et al. 1999), and multi-person coordination (Veeneman & Lee 2000), its ability to truly predict complex, multi degree of freedom, 'real' behaviours has yet to be tested. However, dynamical pattern theory as a whole has been the driving force between a 'paradigm shift' within the motor control community. Researchers now must at least consider the dynamical correlates of information processing constructs such as the *executive* and *motor program*; proponents of new approaches must similarly be able to account for phase transitions and more general non-nonlinear behaviours.

2.8 Conclusions

Table 2.8 summarizes the different types of models discussed to this point; this demonstrates how each model manages only to describe a subset of human behaviour and have limiting model assumptions.

1) Current modeling approaches only indirectly address the issues involved with variability in motor performance. Consider the HKB-model and the stochastic optimized submovement models; both models explicitly acknowledge the indeterminacy of movement. These hypotheses however both consider the indeterminacy to be the result of *noise* in the system which "reduces" performance. This approach is taken to the extreme in the HKB model, simply 'adding a noise term' to a deterministic process. However, viewing variability as noise is at odds with recent indications that motor variability has 'non-white' elements: variability should instead be an "intrinsic" rather than "explicit" component of action models. Another shortfall comes from the methods for forming the internal properties of models. With the exception of the VITE model (and derivatives), and to a limited extent the minimization models, researchers construct the parameters of their models based on *a priori* hypotheses regarding structure and function. For example, Rosenbaum and colleagues make a variety of assumptions regarding rule priority (cost).

Table 2.1: Summary of current models in motor control

Model	Properties			Calculations	Storage	Basis
Equilibrium Point	2-D	deterministic + noise	open loop	moderate	low	muscle mechanics
HKB	1-D	deterministic + noise		moderate	low	synergetics
Knowledge	2-D	deterministic	open loop	moderate	high	psychology
VITE	1-D	deterministic	closed loop	moderate	moderate	neural networks
Minimum Jerk	2-D	deterministic	open loop	high	low	robotics
Minimum Torque	2-D	deterministic	open loop	v. high	low	robotics
Iterative Corrections	1-D	deterministic	closed loop	moderate	low	visual control
Optimized Submovement	1-D	stochastic†	closed loop	moderate	low	visual control
Kinematic Theory	2-D	stochastic‡	open loop	moderate	low	robotics

† The stochastic nature of model is entirely based on variability as an expression of neuromotor noise

‡ Although the model is often interpreted in a deterministic fashion, the structure of the model is probability based

2) With few exceptions (e.g., VITE), current models do not consider the process of perceptual motor learning; skill acquisition is in many cases assumed to be pre-encoded. Moreover, while some models do provide implications for learning (e.g., the HKB-model), even providing examples of learning adaptation (e.g., Schöner, Zanone & Kelso 1992), no mechanism is provided for such changes. However, as we will see, learning plays a key role in modeling: for parameter estimation, for skill generalization, and for performance evaluation.

An alternative view to assuming model structure is to proceed based entirely on the observed data, acquiring the internal form from data-based model update.[‡] While some current approaches employ neural network learning methods (e.g., Bullock et al. 1998, adaptive resonance theory and backpropogation), others similarly draw from the machine-learning literature which provides aspects unavailable or impractical to neural networks.

2.9 Summary

The previous discussion outlined a number of the more influential models of human motor control, which describe a wide range of behaviours. However, as was suggested earlier, the goal of this dissertation is to develop a 'general' framework for studying perception-action systems based on perceptual learning, implicit variability, and the ability of fundamental principles to apply to many tasks. Reflecting on the successes and failures of the discussed models there are some key features that should be included in any new approach:

1: Use current knowledge regarding system organization (both anatomical and process related) within the model structure;

2: Place reasonable[§] cognitive and computational demands on the system

3: Capture the inherent dynamic variablity of performance within the process model;

[‡]This is the extreme of this viewpoint; the researcher may well wish to constrain the model update procedure based on hypotheses.

[§]"a minimally intelligent executive intervening minimally" (Kugler et al. 1980)

4: Acquire model parameters via learning (supervised or unsupervised) to describe the dynamics of the target;and

5: Apply to a wide variety of tasks and levels of analysis.

Bearing in mind these points, a class of modeling procedures based on a probabilistic (Bayesian) viewpoint are pursued.

Chapter 3

A Computational Intelligence Perspective

The problem of understanding and representing multivariate signals has a long tradition within the computing/engineering literatures. From an Artificial Intelligence (AI) standpoint, brain/motor function is contrived as some kind of computation, and the ability to simulate human performance provides insight into "biological algorithms". One problem encountered by early AI research was based on the assumption that information was processed serially and symbolically, similar to information processing approaches in psychology. Using this serial premise, computer processing speeds were unable to approach even simple reaction times measured in humans (Stacey 1998). However, considering the system as massively parallel and *open* may alleviate some of these shortfalls.

In 1985, Carl Hewitt published "The Challenge of Open Systems", which suggests the human system is subject to communications and constraints from the outside world and is thus characterized by (Hewitt 1985a, Hewitt 1985b):

Concurrency. Open systems are composed of numerous interacting components which must process information concurrently.

Asynchrony. New information may enter the system at any time, preventing synchronous interaction with the outside world. The components of the system are separated by distances (physical or philosophical) preventing them from synchronizing.

Decentralized Control. Any central decision-maker comprises a substantive bottleneck. Due to the asynchrony and unreliability of communication, the homunculus would never have access to complete and timely system-state information. Thus the system requires distributed control, with decisions being made local to each process.

Inconsistent Information. Due to the possibility of inaccuracies of information from within or without the system, decisions must be made by the components of the system.

Arm's-length Relationships. The internal state and/or action of one agent or component may be unknown to other agents.

Continuous Operation. Open systems must be designed to resist system failure due to local (agent) events.

Interestingly, the 'open' approach forwarded by Hewitt criticizing many of the serial and localized function assumptions of earlier AI views, is reminiscent of the recent 'levels of analysis debate' within the motor control literature. In this dialog, the concept of *self-organization* enables highly organized behaviour to arise from the interaction between unknown microscopic dynamics (e.g., Kelso 1984, Kugler et al. 1980, Kugler et al. 1982, Kugler & Turvey 1987). Moreover, in abstract systems the emergent behaviour is only meaningful from a macroscopic viewpoint; at a microscopic level the transitions are indistinguishable from randomness (Stacey 1998).

In order to represent the emergent behaviour of the human motor control system dis-

cussed previously, it is reasonable to pose the problem in terms of a dynamical system. Such a system is constrained by a set of initial conditions and passes through a succession of states, similar to formulations discussed earlier for coupled oscillator models (Kelso 1984, Haken et al. 1985)*. A dramatic benefit of such a dynamical systems approach is that conclusions are based on the application of *evolutionary* rules or transformations, having no need for a controller (c.f., Hewitt 1985a, Hewitt 1985b).

For the purpose of the present work, two methods of distributed computation are discussed with regard to understanding motor behaviour. The first, neural networks, are cursorily discussed, covering the basic computational methods needed to elucidate the underpinnings, benefits, and weaknesses of some current motor control models†. The primary focus of this chapter however is the implementation and implications of Hidden Markov models(HMM). This structure for dynamical system estimation and representation is presented in a thorough manner to enable further discussion and application for motor control modeling. In later chapters HMMs are extended within a Bayesian belief structure for the modeling of more complex behaviour.

3.1 Basic Methods in Artificial Neural Networks

Artificial Neural Networks (ANNs) are a recognized method for learning real, discrete, and vector valued functions (Mitchell 1997), with applications ranging across many disciplines. Moreover, ANNs are able to generate outputs quickly despite their relatively long training times; outputs are produced by the arithmetic combination of weighted nodes without requiring any run-time estimation. These networks have the added benefit, for those interested in biological systems, of a structure that may be used as a

*An attractor is a state, or state vector, to which the system may evolve. The *basin of attraction* of an attractor is formed by the set of initial conditions which might give rise to the system terminating in that attractor.

†A sub-category of neural networks (cascade networks) is applied to the calibration of collected data (see Appendix A).

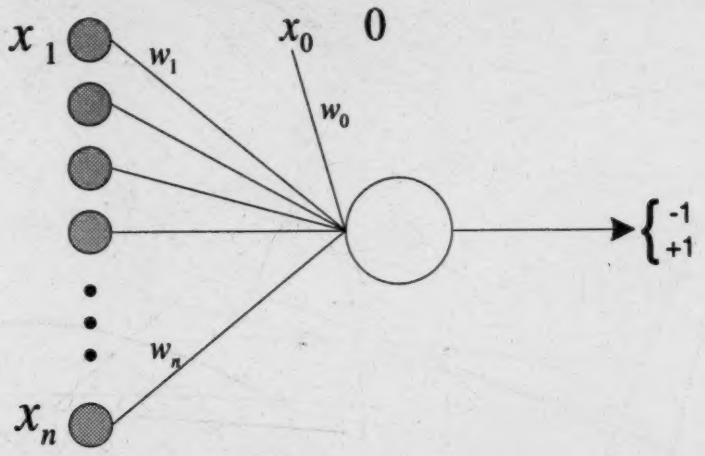

Figure 3.1: Graphical schematic of a simple perceptron.

metaphor to cellular formations (Hebbian networks).

3.1.1 Perceptrons

The simplest neural network, the perceptron, transforms a vector of real numbers into an output of 1 or -1 (yes or no). In general, the output may be defined by the linear combination of the weighted inputs ($w_0 + \ldots$), where the precise weights are learned) in relation to some threshold (Equation 3.1). The input weightings, w_i's, are real valued and $-w_0$ is a learned threshold value such that:

$$o(x_0, \ldots, x_n) = \begin{cases} -1 & \text{if} \quad w_o + w_1 x1 + w_2 x_2 + \ldots + w_n xn > 0 \\ +1 & \text{if} \quad \text{otherwise} \end{cases} \tag{3.1}$$

Initially, the weights of such a network are acquired iteratively based upon an initial set of random values. The initial perceptron is applied to each of a set of training examples; each time an incorrect classification is made, an update is performed according to a training rule which revises the w_i associated with each x_i (see Equation 3.3). One

example, which requires training examples to be linearly separable[‡], is the perceptron rule:

$$\tilde{w}_i = w_i + \Delta w_i \tag{3.2}$$

$$\Delta w_i = \eta(t - o)x_i \tag{3.3}$$

where the model update value (Δw_i) is determined by the difference between the observed and target output (o and t respectively), scaled by learning rate, η(positive constant; e.g., 0.1).

Unlike the perceptron rule which updates the weight coefficients based on the thresholded $(+1, -1)$ value, rules based on gradient descent operate on an unthresholded value – relieving the necessity of linear separability. Here the weights are adjusted in the direction of steepest descent in all dimensions of $\vec{w}$ (see Equation 3.4). In standard implementations of gradient descent, the error surface may be determined either on an example to example basis or over all training examples before updating is performed (Mitchell 1997). However, both methods can be considered to be following the steepest descent in weight space provided that step size (learning rate) is adequately small.

$$E_p = \frac{1}{2}\sum_j (t_j - o_j)^2 \tag{3.4}$$

3.1.2 Backpropagation

While the the previous examples were convenient for understanding ANN structure, most common applications of ANNs require the mapping of input data onto multiple outputs (see Figure 3.2). In order accomplish such transformation, additional layers, generally composed of nodes containing so-called squashing functions (e.g., logistic function), are added to the network. Thus the task of the learning algorithm is to

[‡]Classification may be made in hypothesis space by a straight line.

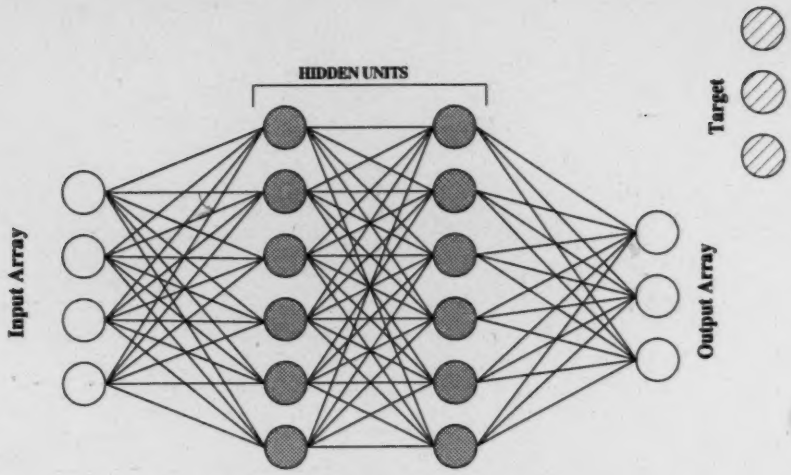

Figure 3.2: A Multi-Layer Neural Network

search across this large hypothesis space for a global minimum error. Gradient descent is once again applied to a network with a fixed number of units and interconnections between units; this process is called *backpropagation*. Initially, all weights are set to small random numbers (e.g., 0.05 to -0.05). Until the termination state is reached the following processes occurs:

1: Propagate Forward Error

i. input $\vec{x}$ and compute o_u for every output unit in the network

2: Propagate Backward Error

i. for each output unit k, calculate an error term d_k where:

$$d_k \leftarrow o_k(1 - o_k)(t_k - o_k) \tag{3.5}$$

ii. for each hidden unit h, calculate an error term d_h where:

$$d_h \leftarrow o_h(1 - o_h) \left(\sum_{k \epsilon outputs} w_{kh} \delta_k \right) \tag{3.6}$$

iii. update all w_{ji}'s

$$\tilde{w_{ji}} = w_{ji} + \Delta w_{ji} \tag{3.7}$$

$$\Delta w_{ji} = \eta \delta_j x_i \tag{3.8}$$

In general, each example forms a training pair $\vec{x}, \vec{t}$. The network is then applied to the example vectors. The output error of the network is then estimated and a gradient calculated; dependent on the chosen method, this procedure may be done after each example, (*stochastic approximation*), or after all examples. Weights are updated in the direction of maximum descent; this step is continued, often thousands of times, until the termination is criteria is reached. Termination may be determined by a number of methods such as total interation count or error threshold. The choice of termination condition is important as reaching this state too early will reduce performance while reaching this state too late may overtrain the network thus limiting is generalizability (Mitchell 1997).

3.1.3 Recurrent Networks

The previous discussion focused on feedforward networks whereby input data is mapped via a number of hidden units and weights onto output nodes. While the exact structure of the network and the nature of the output may vary significantly, feedforward networks apply only to time-invariant systems. This assumption of temporal stationarity is reasonable for applications such as calibration (see Appendix A); there is no expectation for the electromagnetic properties of the collection environment to change over time. However, further discussion here will focus specifically on biological systems where this assumption would be inappropriate.

The alternative to feedforward networks is termed a *recurrent* neural network, where the output data from units at time, t, act as inputs to other units at $t + 1$. In this

representation, a new unit, b, is added to the hidden layer, recording the current level of activation across the hidden units. Similarly, a new input $c(t)$ is added to the input (see Figure 3.3); this unit is a copy of b from the previous time. These units represent the dynamical nature of the system, allowing the internal properties of the system previously to impact the output of the current model. The difficulty of training such a network can be seen by the dependence of $o(t+1)$ on both $x(t)$ as well as $c(t)$ which is likewise dependent on network events at $t-1$. While training methodologies for recurrent networks are still an item of active research (e.g., Jordan 1996, Mozer 1995), in general backpropagation may be applied over a temporally expanded network (see Figure 3.3). Using this technique, the final weights are taken as the mean value of the corresponding coefficients for all time T. Recurrent networks implementing time-dependent back-propagation or similar methods have demonstrated utility for modeling human motor control systems due to their ability to identify (and hence control) non-linear dynamical systems (Narendra 1995). In particular, they may be applied to the estimation of many system configurations including feedforward and internal model (feedback) variants (Narendra 1995).

3.1.4 Applications in Motor Control

3.1.4.1 *VITE* Revisited

In the *VITE* model discussed earlier, Bullock and colleagues implemented a neural network for the examination of reaching to grasp. Unlike all of the supervised network systems discussed so far, *VITE* and its extensions (DIRECT, FLETE, AVITE) combine a series of unsupervised neural architectures. This is accomplished by applying local learning laws to initial random neuronal weights, thus permitting self-organized transformations and classifications. Autonomous learning is accomplished through random, repetitive, internally generated action-perception-cycles (see Figure 3.4). Specifically, this 'motor babbling' is achieved via an Endogenous Random Generator (ERG) whose

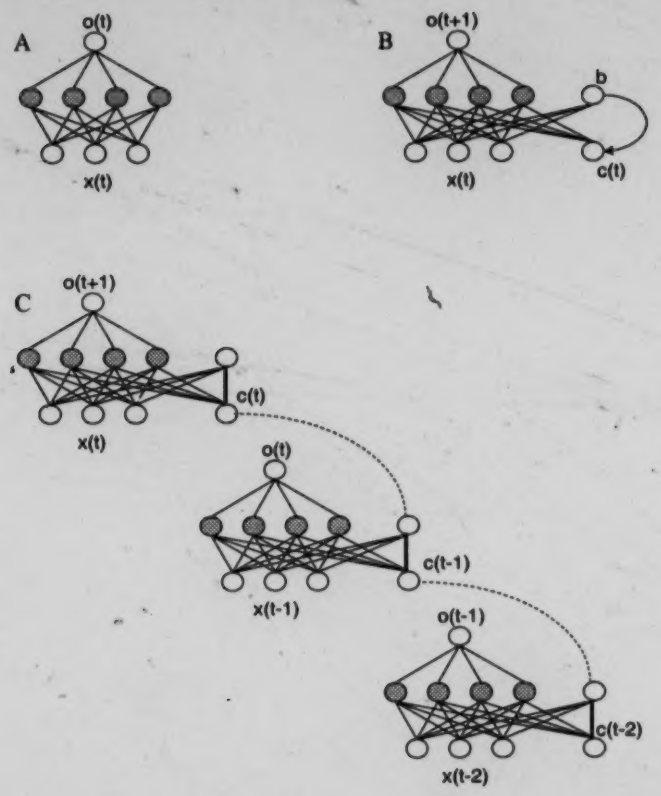

Figure 3.3: Schematic of an 'unfolded' representation of a recurrent neural network.

output triggers motor commands (Gaudiano & Grossberg 1991). Looking specifically at the AVITE implementation of this emergent learning, several features should be noted. In the model's initial formation the ERG output is used to generate random movements (Figure 3.4a); this initialization elegantly avoids the *novelty problem* discussed earlier for motor programming solutions to movement. The motor output is subsequently used as an 'intended' movement outcome (TPC). Given a paired input and output, the spatial-motor command filter may be acquired by multiple iterations, thus acquiring the transformation by randomly sampling the workspace, independently

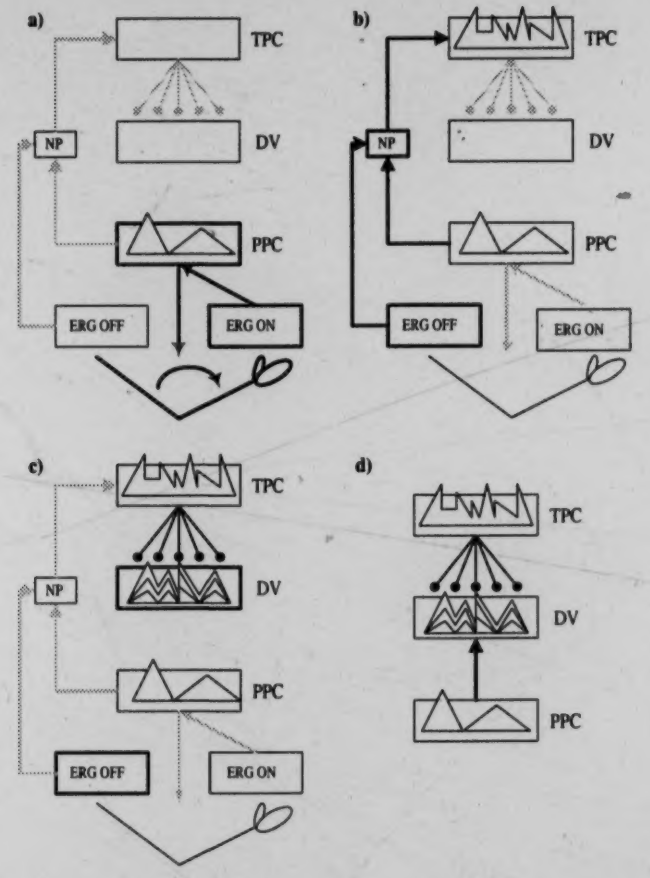

Figure 3.4: Block diagram of AVITE circuit during autonomous learning – "motor babbling". (a) The ERG channel output is input as a difference vector to the PPC (present position command) node. A random arm motion is generated. (b) At ERG offset, the arm ceases movement and the ERG OFF gate incurs the now print command (NP); a copy of the current PPC is induced as a TPC (target position) command. (c) The TPC command is filtered via the 'to be acquired' spatial-motor transformation and compared with the input PPC command at the difference vector node (DV). (d) The learning law adjusts the $TPC \longrightarrow DV$ synapses such that $DV = 0$.

(From Gaudiano & Grossberg, 1991)

of the physical parameters required for muscle/joint preparation.[§] In addition to the features this provides for initial skill acquisition, added benefits are gained by the generation of an internal error measure – DV. The DV signal can be used in an on-line

[§]The exact physical implications of any acquired transformation could be acquired by increasing the extent of the model to include node for force-length and associated programming variables.

regulation of both performance and learning. Specifically, during visual regulation of a reaching movement, the DV error value reflects the current pointing error, allowing a corrective movement to be made (c.f., optimized submovement model of control; Chapter 2). The DV signal in combination with a reinstatiation of the learning loop (NP node activation, see Figure 3.4) enable the current movement error to be implemented in the spatial-motor transformation, thus reducing future movement errors.

3.1.4.2 Minimization Procedures

In an earlier discussion, several example models utilized minimization of a control parameter as a pivotal heuristic. (e.g., Flash & Hogan 1985, Rosenbaum et al. 1993, Rosenbaum et al. 1999, Uno et al. 1989). While the performance of this approach has demonstrated impressive descriptive power, the process of optimization is far from trivial. Although, other alternatives exist to accomplish the minimization, the neural network methods will be discussed due to its biological plausibility (Kawato 1996).

In general, any biological neural network implementation must perform under the computational constraints of the target system – in this case the human action-perception system. Therefore a minimization network for torque, jerk, or *cost* (Rosenbaum et al. 1993, Rosenbaum et al. 1999) must:

1) Acquire, by training, an internal representation of the physical system; and

2) Simultaneously solve (a) trajectory formation (b) coordinate transformation and (c) system control problems.

There are many possible implementations of appropriately constrained, neural network based, minimization. One relatively straightforward approach had Massone and colleagues (Massone & Bizzi 1989) train a recurrent neural network to generate a minimum-jerk trajectory for a 2-link and six-muscle arm system. During supervised

learning, the network was provided with a 'perfect' teaching set of the desired 'minimized' trajectory. Thus demonstrating that, at least under ideal circumstances, a simple supervised recurrent learner can acquire and store the necessary transformations. A more involved method of optimization utilizes a cascade neural network to dynamically adapt to environmental features such as obstacles and desired via-points. For example, Kawato et al. (Kawato 1990) represented the dynamic properties of a system spatially in the network. As such, the 'motor command' nodes are fully connected to those at previous epochs, allowing a forward model of the network to be acquired through backpropagation and motor commands to 'control' the object in accord with minimum torque change.

3.1.4.3 Neural Networks - Parting Thoughts

Artificial neural networks present many benefits beyond the superficial similarities between their structure and that of biological nervous systems. In particular, ANNs are able to acquire their internal structure based on known input-output pairings (supervised) or by applying local learning rules to find the appropriate classifications/transformations. Some formulations are also able to represent motor control features such as output variability (Probabilistic Neural Networks) while retaining impressive robustness and some degree of generalization. However, there are some noticeable shortcomings of ANNs. Specifically when the networks are (1) highly parallel/connected, (2) modular, and (3) recurrent (increasing with order), optimization and updated procedures have shown to be susceptible to finding sub-optimal solutions (Caelli, Guan & Wen 1999).[¶]

[¶]There is a tendency for methods such as recurrent backpropagation to get 'stuck' in local minima of the error landscape (Caelli et al. 1999).

3.2 First-Order Markov Chains

An alternative to the network/lattice architecture of the ANN is the 1^{st}-order Markov chain. While the central concerns regarding ANN use in biological systems are situations where behaviour is dynamical and the known anatomy appears modular; first-order Markov chains are inherently dynamical and may be represented within modular arrangements.

One of the central assumptions of all discrete 1^{st}-order Markov chains is that a system may be described as a series of N distinct states, q. The system may undergo change, between states, at discrete time intervals ($t = 1, 2, ...,$). The initial likelihood of a state is expressed by π_i while the transitions between states are governed by a set of state transition *probabilities* a_{ij} such that:

$$a_{ij} = p[q_t = j | q_{t-1} = i] \tag{3.9}$$

where

$$a_{ij} \geq 0 \qquad \forall j, i \tag{3.10}$$

$$\sum_{j=1}^{N} a_{ij} = 1 \qquad \forall i \tag{3.11}$$

Thus, the probability of the system being in a given state, q_t, is only dependent on the previous state, q_{t-1}. As a simple example, consider a model for weather prediction (Rabiner & Juang 1993) that is composed of three states: rainy (R), sunny (S), cloudy (C). The transition probabilities (a_{ij}) are represented by the stationary (time invariant) transition matrix A:

$$A = \begin{array}{c|ccc} & (R) & (S) & (C) \\ (R) & 0.60 & 0.20 & 0.20 \\ (S) & 0.33 & 0.33 & 0.33 \\ (C) & 0.30 & 0.50 & 0.20 \end{array} \tag{3.12}$$

Based on A we may evaluate the probability of an observation sequence or the likelihood

of a state being the next observation. For example, the probability of a series of days following the sequence rainy, rainy, sunny, rainy, and cloudy (replacing 1 for rainy, 2 for sunny, 3 for cloudy):

$$\begin{aligned} P(\mathbf{O}|Model) &= p[1,1,2,1,3|Model] \\ &= p[1]p[1|1]p[2|1]p[1|2]p[3|1] && (3.13) \\ &= \pi_1 \cdot a_{11} \cdot a_{21} \cdot a_{12} \cdot a_{31} && (3.14) \\ &= (0.33)(0.60)(0.2)(0.33)(0.2) && (3.15) \\ &= 2.61 \times 10^{-3} && (3.16) \end{aligned}$$

where π_1 is the initial probability of state 1 (rainy) given no history.

Although somewhat simplistic in presentation, Markov chains are a valuable tool in representing non-deterministic processes (like weather). By representing the system in this *stochastic* manner, variability in output, given a constant input, is captured by the structure of the model. Similarly, the model may be adjusted to reflect the magnitude of indeterminacy between events. Although Markov chains are adept at representing relatively simple finite state systems, given the goal of representing human behaviour here, they are not easily generalized to more complex systems – specifically systems where the underlying process is unobservable or unknown.

3.3 Hidden Markov Models

In addition to Markov chains relying on observable events they also have an inherent assumption of stationarity – the relation between state t and $t+1$ remains constant over time. Alternately, hidden Markov models (HMM) include a second level of probabilities (i.e., doubly stochastic), allowing HMM's to represent substantially more complicated

systems. The thought experiment presented by Rabiner (1993) of coin tossing provides an excellent summary of the properties of HMM and the problems encountered during their application.

In this experiment, all the subject is told is the output of the current toss (i.e., Heads or Tails). The situation gets more complicated however, because the experimenter may have more than one type of coin and does not say which coin was flipped to result in the observation of heads or tails (the selected coin is hidden). Given only the sequence of observations ($O = o_1 o_2 o_3 \dots o_T$) the subject must predict what the next observation will be. In order to succeed, the following must be known (or estimated):

1. How many coins (states) are there?
2. Given a coin, what are the chances of it producing heads or tails (i.e., what is the coin's bias)?
3. What is the probability of choosing a given coin for the next toss (i.e., what is the relation between coins)?

3.3.1 Definitions

Prior to a detailed discussion of HMM algorithms and architecture, it is useful to define a number of terms and variables which are used throughout the remainder of the chapter. In order to retain meaning to the terms, definitions are given with regard to an HMM representing the previous coin example.

T = the length of the observation sequence

N = the number of states in the model (e.g., the number of coins). Similarly to the Markov chain, states are labeled $(1, 2, \dots, N)$ and the state at time t is q_t

M = the number of possible observations per state (e.g., heads(H) or tails(T))

A = the state-transition probability distribution a_{ij} (the same as in the Markov chain).

For example, given 3 coins:

$$A = \begin{array}{c|ccc} & (1) & (2) & (3) \\ (1) & 0.10 & 0.10 & 0.80 \\ (2) & 0.33 & 0.33 & 0.33 \\ (3) & 0.25 & 0.25 & 0.50 \end{array} \tag{3.17}$$

B = the state-dependent observation probability distribution $b_j(k)$, where:

$$b_j(k) = p[\mathbf{o_t} = \mathbf{v_k} | q_t = j], \quad 1 \leq k \leq M \tag{3.18}$$

$$B = \begin{array}{c|cc} & (Heads) & (Tails) \\ (Coin1) & 0.50 & 0.50 \\ (Coin2) & 0.60 & 0.40 \\ (Coin3) & 0.10 & 0.90 \end{array} \tag{3.19}$$

defines the probability of an observation (H,T) given a state (coin: 1,2,3)

π_i = the initial state distribution (again, the same as in the Markov chain) representing the distribution of states as initial values.

$$\pi_1 = 0.33 \quad \pi_2 = 0.33 \quad \pi_3 = 0.33 \tag{3.20}$$

λ = the model, containing a complete set of parameters ($\lambda = (A, B, \pi)$).

3.3.2 Classical Estimation Problems in HMMs

There are three types of problems in model estimation for HMMs. They are:

1. Given a sequence of observations $\mathbf{O} = (\mathbf{o_1 o_2} \ldots \mathbf{o_T})$ and a complete model λ, compute the likelihood of the observation sequence given the model $p(\mathbf{O}|\lambda)$.

2. Given the model λ and an observation sequence $\mathbf{O}$, select the most probable state sequence $(q_1 q_2 \ldots q_T)$.

3. Adjust the model, $\lambda = (A, B, \pi)$, to maximize the likelihood of the observed sequence, $p(\mathbf{O}|\lambda)$.

Just as with neural networks, HMMs acquire their internal structure from recursive update algorithms. While ANNs learn via error propagation through the network to update the weights, HMMs base update on a process termed "Expectation Maximization" (EM). In this process the likelihood of each datapoint of the observed system is estimated given a model (initially random); the model is updated based on these probabilities via an iterative algorithm known as Baum Welch. The underlying "hidden" process of the model can be estimated by a variant of EM which selects the most likely sequence of states given the observations and the new model – the Viterbi algorithm.

While little attempt is made in this chapter to frame the algorithms within a motor control setting, a brief preview of Chapter 5 will serve the reader well while considering the methods of optimizing. Consider the simple control system of one-dimensional pointing to a object. Let us define the position of end-effector of the arm relative to a target $(o_1 \ldots o_T)$ and hypothesize a first-order controller (described by n-states and their transitions). The goal of our update procedures will then be to estimate the statistics (transition matrix, $\mathbf{A}, \mathbf{B}$) directing the control states based on the observed arm position.

3.3.3 The Viterbi and Baum-Welch Procedures

Prior to defining these algorithms, formal definitions of the forward operator $\alpha_t(i)$ and the backward operator $\beta_t(i)$ are discussed as they represent the fundamental search method incorporated in both the Viterbi and Baum Welch algorithms.

3.3.3.1 The Forward Operator

The forward variable $\alpha_t(i)$ is defined as the probability of being in state i at time t and observing a sequence from time 1 to time t ($o_1 o_2 \ldots o_t$) given the model λ (Equation 3.21). This is calculated for all states M and, using the coin-tossing example, gives the probability of the observation sequence $O = (h, h, h, t, t, h, t)$. This is accomplished by evaluating $\alpha_t(i)$ for each observation in turn based on the previous $\alpha_t(i)$ for all states.

$$\alpha_t(i) = p(o_1 o_2, \ldots o_t, q_t = i|\lambda) \tag{3.21}$$

A solution may be generated for $\alpha_t(i)$ by three general steps. Initially, for $t(1)$, $\alpha_t(i)$ is set for all states i based on the initial probability of state i (π_i) and the probability of the observation (o_1 from the sequence) given state i (Equation 3.22). The induction step (see Figure 3.5, Equation 3.23) demonstrates how state j at time $t+1$ may be reached from all N states at time t. The chance of being at state j at time $t+1$ is based on the summation of the probability of being in all N states at time t ($\alpha_t(i)$ for all states),the state transition probabilities (a_{ij},matrix A), and the probability of observing o_i given state j. This proceeds for all j at time $t+1$ and for the entire observation sequence until time T (Equation 3.24).

Initialization

$$\alpha_1(i) = \pi_i b_i(o_1) \qquad 1 \leq i \leq N \tag{3.22}$$

Induction

$$\alpha_{t+1}(j) = \left[\sum_{i=1}^{N} \alpha_t(i) a_{ij}\right] b_j(o_{t+1}) \qquad \begin{array}{l} 1 \leq t \leq T-1 \\ 1 \leq j \leq N \end{array} \tag{3.23}$$

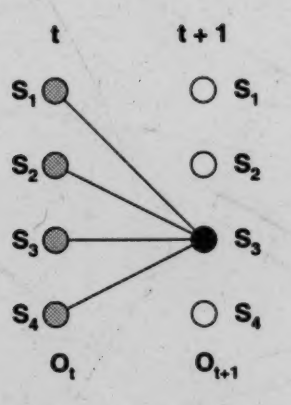

Figure 3.5: Graphical representation of the forward operator. Iterations begin at time t_1 and search forward to T.

Termination

$$p(O|\lambda) = \sum_{i=1}^{N} \alpha_T(i) \tag{3.24}$$

Based on this lattice-like procedure, all possible state sequences that will result in the given O and state are tested, regardless of the length of the sequence.

3.3.3.2 The Backward Operator

Operationally, the backward variable $\beta_t(i)$ is very similar to the forward operator and may be defined by the probability of the partial observation sequence $(o_{t+1}o_{t+2}\ldots o_T)$ given state i at time t and the model λ (Equation 3.25). This operator seeks to find the probability of an observation and state based the points following it in time; the search proceeds by assessing the probability of state, given O, based on the $\beta_t(i)$ for the all preceding states.

$$\beta_t(i) = p(o_{t+1}o_{t+2}\ldots o_t|q_t = i, \lambda) \tag{3.25}$$

Initialization

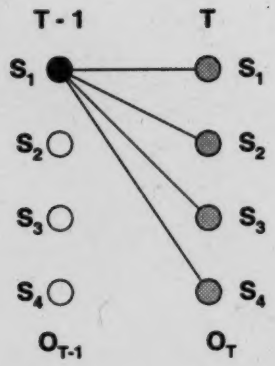

Figure 3.6: Graphical representation of the backward operator. Iterations begin at time T (end of time series) and proceed backward to t_1

$$\beta_T(i) = 1 \qquad 1 \leq i \leq N \tag{3.26}$$

Induction

$$\beta_t(i) = \sum_{j=1}^{N} a_{ij} b_j(o_{t+1}) \beta_{t+1}(j) \tag{3.27}$$
$$t = T-1, T-2, \ldots, 1 \quad 1 \leq i \leq N$$

A value of 1 is assigned to βT for all states and a lattice progression follows, similar to the forward procedure (see Figure 3.6, Equations 3.26 and 3.27).‖

3.3.3.3 The Viterbi Algorithm

Sometimes considered a special case of the forward operator, the goal of Viterbi search is to find the best state sequence for a string of observations, given λ.** Again, from the coin example, what is the most probable sequence of coins (states) to yield the

‖Note the assignment of $t = T-1, T-2, \ldots$ for the progression so that the algorithm proceeds toward t from T.

**Returning to the pointing example the goal here would be to estimate the sequence of *some hypothetical* controller given a model and a sequence of observed arm positions.

observation sequence, $O = (h, h, h, t, t, h, t)$. For each time step ($t$ to $t+1$) a search is performed to find the most likely state ($a_{ij}bj(o_{t+1})$) given the observation (o_{t+1}, see Figure 3.7). In a similar fashion to the forward operator, Viterbi is initialized using π and b_i, however an additional array (ψ_i) is also set to allow the recording of the 'found' sequence (Equation 3.28 and 3.29) .

Initialization

$$\delta_1(i) = \pi_i b_i(o_1) \quad 1 \leq i \leq N \tag{3.28}$$
$$\psi_1(i) = 0$$

Recursion

$$\delta_t(j) = \max_{1 \leq i \leq N}[\delta_{t-1}(i)a_{ij}]b_j(o_t) \quad 2 \leq t \leq T, 1 \leq j \leq N \tag{3.29}$$
$$\psi_t(j) = \arg\max_{1 \leq i \leq N}[\delta_{t-1}(i)a_{ij}] \quad 2 \leq t \leq T, 1 \leq j \leq N$$

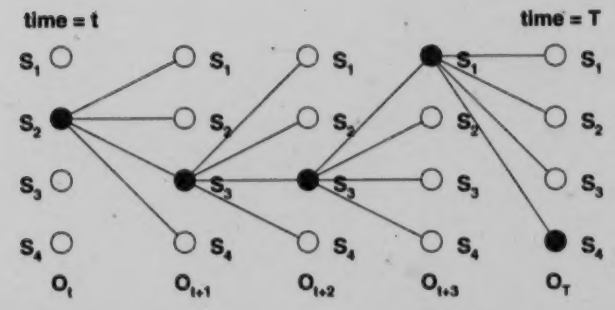

Figure 3.7: Graphical representation of the Viterbi search procedure. The optimal state sequence is estimated based on an *optimal* search. Here the most likely way to get to each state at each time interval is recorded (see Equation 3.29, ψ_t); joint state (a_{ij}) and state-given-observation (b_o) likelihoods are evaluated 'forward', building a list of possible states given the observations. The state trellis is then used to define the most likely state sequence. Only the state transitions are considered during this recursion step to select the final state sequence from the list.

During recursion, the procedure continues to select the most probable state-given-observation for all time steps with the resulting state sequence being recorded within

$\psi_t(j)$. Termination is reached at the end of the observation sequence (T). During model establishment, the 'fit' of the observation sequence is then evaluated both by examination of the overall likelihood of the sequence and by the number of 'replacements' which would be necessary to produce the correct state sequence (see Section 3.3.4). Finally, the optimal path through the "state trellis" is determined by back tracking to find the sequence of states with the highest probability of occurring.

3.3.3.4 The Baum-Welch Algorithm

The purpose of the Baum-Welch Algorithm is to update the current model with respect to the observation sequence, thus increasing the success measures following a Viterbi search. In order to accomplish this, two values are required:

$$\xi_t(i,j) = p(q_t = i, q_{t+1} = j|O, \lambda) \tag{3.30}$$

that is, the joint probability of being in state i at time t and state j at time $t+1$. As one might predict, the forward and backward operators are used to converge on this solution (see Figure 3.8, Equation 3.31). Specifically, the forward operator is run until time t while and the backward operator from T until $t+1$.

$$\xi_t(i,j) = \frac{\alpha_t(i)a_{ij}b_j(o_{t+1})\beta_{t+1}(j)}{\sum_{i=1}^{N}\sum_{j=1}^{N}\alpha_t(i)a_{ij}b_j(o_{t+1})\beta_{t+1}(j)} \tag{3.31}$$

The probability of being in state i at time t, $\gamma_t(i)$, is also required. This quantity is calculated from summing all $\xi_t(i,j)$ across all states j (Equation 3.32).

$$\gamma_t(i) = \sum_{j=1}^{N}\xi_t(i,j) \tag{3.32}$$

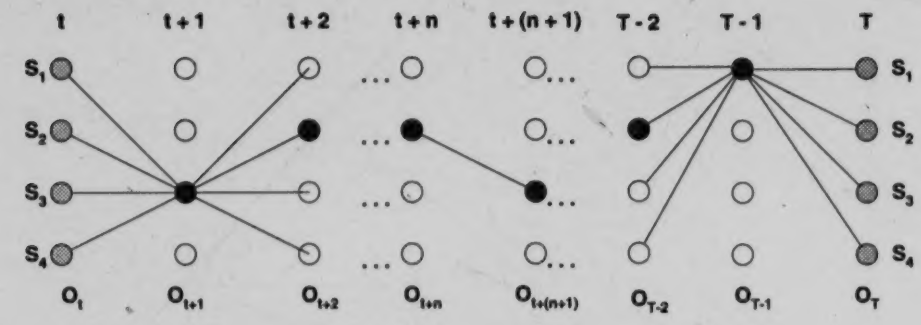

Figure 3.8: Graphical representation of the Baum-Welch model update algorithm. An HMM specific form of expectation maximization, the forward and backward operator are used to find the likelihood of state transition ij.

Based on $\xi_t(i,j)$ and $\gamma t(i)$, the model λ may be re-estimated such that:

$$\begin{aligned}
\hat{\pi}_j &= \text{expected number of occurrences of state } i \text{ at time } t(1)(\gamma_1(i)) \qquad (3.33)\\
\hat{a}_{ij} &= \frac{\text{number of transitions expected from state } i \text{ to } j}{\text{number of transitions expected from state } i}\\
\hat{b}_j(k) &= \frac{\text{expected number of occurrences of state } j \text{ concurrent with observing symbol } k}{\text{expected number of occurrences of state } j}
\end{aligned}$$

where

$$\begin{aligned}
\sum_{t=1}^{T-1} \gamma_t(i) &= \text{expected number of transitions from state } i \text{ given } O \qquad (3.34)\\
\sum_{t=1}^{T-1} \xi_t(i,j) &= \text{expected number of transitions from state } i \text{ to state } j \text{ given } O
\end{aligned}$$

Thus new parameter values for the model $(\hat{\pi}, \hat{A}, \hat{B})$ are generated based on the observation sequence. Following model update, the Viterbi algorithm may be run again to determine the degree to which the update has improved performance of the model or, as will be the case for the present discussion, the model may be used to generate sequences which are then evaluated against the training data.

3.3.4 Assessing HMM Performance

Although the Viterbi algorithm is often used in "recognition" applications to determine the most likely model – thus recognizing the sequence based on the probability of the sequence – Monte Carlo sampling combined with a Hamming distance (see Section 3.3.4.3) measure can be used to generate a *new* sequence and assess the degree of agreement with a measured observation sequence.*

3.3.4.1 Unconstrained Monte Carlo Sampling

The "unconstrained" method reflects a 'full' sampling of both the state and observation distributions to determine the new observation value. Specifically, at each time-point, a new state is randomly selected from the distribution of all possible states given the current state (Equation 3.35); an observation is then selected from the distribution of possible observations given the selected state (Equation 3.36).

$$choose(s_i) \propto p(s_i|s_k) \tag{3.35}$$

$$choose(o_j) \propto p(o_j|s_i) \tag{3.36}$$

where s_i is the current state and s_k is the state at the previous epoch. The observation (o_j) at the present interval is thus dependent on the current state selection s_i.
This method will be used extensively in the following examinations as it represents the most likely utilization of an HMM as a "motor program". Thus a command, in the form of $\lambda\ (A, B, \pi)$, could be selected and initiated centrally and sampled locally to the joint/subsystem.

*Once again considering the arm control will give some degree of concreteness. Here one would generate a new trajectory and perform a point-wise comparison with an observed pointing trial.

3.3.4.2 Constrained MonteCarlo Sampling

As opposed to the previous method, "constrained" MonteCarlo sampling does not select from the state distributions. Instead, the Viterbi generated *optimal* sequence is used to restrict the search; thus, given the Viterbi state (x) at each epoch, randomly select from all observations such that:

$$choose(o_i) \propto p(o_i|x) \tag{3.37}$$

Although constrained sampling is the preferred tool within the machine learning community for prediction – invariably returning observation sequences with higher correspondence to recorded values than the unconstrained method(i.e., Hamming distance) – this method does not seem feasible within biological control. The Viterbi selection would effectively require the observation of the movement prior to the planning of action initially. While some might suggest this is reminiscent of planning, it would also require the central nervous system to perform a beam search across the sequence. This high computational load was a criticism of several earlier models as well as being in contradiction to the goal of having a "minimal" central controller.

3.3.4.3 Hamming Distance

Regardless of the generation method, the predicted observation sequence (O) may then be compared at each epoch with a criterion value (measured sequence):

$$\varsigma(\hat{o}, o) = \sum_{t}^{T} I_{o(t)}(\hat{o}(t))/T \tag{3.38}$$

where

$$I_{o(t)}(\hat{o}(t)) = \begin{cases} 1 & \text{iff } \hat{o}(t) \equiv o(t) \\ 0 & \text{otherwise} \end{cases} \tag{3.39}$$

Thus, if the observed matches the predicted value, a score of "1" is assigned, otherwise "0" (Equation 3.39). This binary score is then averaged across the sequence length (T) to generate a score ($\varsigma(\hat{o}, o)$, Equation 3.38). Although this method does represent correspondence, it is unable to account for phase misalignments.[†]

3.3.5 Hidden Markov Models – Closing Thoughts

While hidden Markov models have received considerable study and application within the pattern recognition community (e.g., Rabiner & Juang 1993, Bose & Kuo 1994, Ghahramani 2000, Wilson & Bobick 2001) HMM methods have as yet had few applications to perception-action modeling of human systems. One exception should be noted however (Rimey & Brown 1990). In this report the authors implement a version of HMM, augmented HMMs which allow limited feedback, with the aim of modeling selective attention based on ocular point of regard. Unfortunately, the goal of the model was once again recognition (although limited constrained MonteCarlo trials were performed), however the results do indicate that HMMs in some form may be applied with success to movement production.

[†]Hamming (inverse) distance in this form can be considered a binary comparison of a form very similar to the χ^2 statistic.

Chapter 4

Hidden Markov Models in Action

4.1 Experiment 1: The Simplest Case

As has been demonstrated by the discussion of extant models, the problem of describing and understanding the human control system can be addressed at many levels. Indeed, the researcher may utilize tools ranging from experimental psychology to those derived from principles of synergetics/physics. Similarly, an examination may focus at levels from mechanistic/anatomical to behavioural. The goal of this experiment is to examine the utility of hidden Markov models for studying and representing human actions at a behavioural level. Due to the infancy of this line of research, a relatively simple action, elbow flexion, is modeled. An explicit attempt is made to make minimal assumptions regarding the hidden control process other than those required by HMM structure.

4.1.0.1 Hidden Markov Models for One-dimensional Aiming

As was discussed extensively in Chapter 3, hidden Markov models provide a method for encoding the inherent variability of the motor system as an integral part of a motor

"memory" model. Events are joined in an non-deterministic fashion (e.g., a only sometimes leads to b); from a stimulus-response standpoint, this allows an identical event to cause multiple responses (e.g., a critical value for τ may result in a catching response or a ducking response depending on previous events). Similarly, even given the same intended response, there may be some variability in the actual output. In a fashion similar to neural networks, HMMs also allow the exact mechanism underlying the control to remain hidden, thus allowing the researcher to make no assumptions regarding feedback and other hypothesized events.The goal of the present investigation is to describe/model the action of one-dimensional pointing as a first-order dynamical system. Instead of describing this process with differential equation(s) as is traditional within studies of motor control, hidden Markov models are employed to infer the properties of the 'hidden' controller from the observed limb positions.

Consistent with this formulation of control, the foundation for many joint-centred approaches can be traced to hypotheses suggesting joint control via Tsetlin automata (for a review see Berkinblit et al. 1986).* In their 1986 position paper, Berkinblit and colleagues convincingly assert how a biological control system can utilize a local control structure such as automata (or HMMs); this is taken as a theoretical justification for equilibrium point control. The specifics of the Berkinblit model however have particular relevance for the current model discussion. In particular, these authors take the very strong theoretical stance that joints behave independently, as a function of the limb effector position in relation to the target. Within this supposition, they also proffer a version of the Bernstein problem aptly phrased with regard to reaching behaviour:

> If one takes the view that for each target point ... a table of values of the corresponding joint angles is compiled, then it is necessary to assume that there are many such tables – one for each configuration of the body.
>
> (Berkinblit et al. 1986, pp. 143)

*Tsetlin automata are fundamentally similar to Markov chains. In a Tsetlin processes, a probability density determines the likelihood of transition; however the destination state is predetermined.

Figure 4.1: Annotation frames of an elbow rotation task.

Not unlike Rosenbaum and colleagues' *Knowledge Model*, Berkinblit suggests a relatively simplistic algorithm for ascertaining future postures, based on connected but non-communicating automata, rather than requiring a seemingly infinite look-up table of requisite joint angles; an approach quite consistent with the idea of a "minimally intelligent" controller proposed by Kugler and others (e.g., Kugler et al. 1980, Kugler & Turvey 1987). The approach taken in this study applies the previously discussed HMM in a manner consistent with this distributed approach, situating the bulk of "responsibility" for control at the joint (peripheral) level (Latash 1993, Latash & Gottlieb 1990, Feldman 1986). Specifically, joint control is posed in terms of an advancing hidden Markov model, with observed positions emerging as a result of state transitions; the underlying principles of control are reflected by the model parameters. Examining joint control from this standpoint tests two points: the appropriateness of local joint behaviour as a principle of control and the efficacy of HMMs for representing simple actions.

4.1.1 Methods

4.1.1.1 Participants

Four male subjects volunteered to participate in this investigation (mean age = 26 years). All participants were right handed and naïve to the purpose of the experiment.

4.1.1.2 Task and Procedure

Each subject completed a total of 60 elbow-pointing trials with their right arm: 15 to each of 4 targets (Figure 4.1). On each trial the participant was instructed to flex their elbow, rapidly and accurately, with the goal of aligning their index finger with the designated target. All trials began with the elbow fully extended.

4.1.1.3 Apparatus

Movement data was collected using a Polhemus Ultra-Trak Pro magnetic tracking system at a sampling rate of 120Hz. The collection area was calibrated using a 36 hidden unit (1 layer) cascade neural network (see Appendix A, accuracy < 1 cm), implemented within an in-house data display and analysis package. All incoming data was filtered using a Savitsky-Golay least-squares polynomial filter (4^{th} order, 21 sample window) prior to further reduction. Elbow angle was determined based on magnetic sensor positions(Figure 4.2; red circles). Sensors 1 and 3 were positioned on the lateral aspect of the forearm and upperarm respectively (approximately mid-segment). Sensor 2 was placed over the lateral head of humerus. Full elbow extension was designated as zero degrees flexion. A fourth sensor was placed at the target location to allow calculation of a difference angle between forearm and target.[†] The difference angle was clustered

[†]Estimates of joint angle for the elbow were confirmed based on change in the pitch angle of the sensor 1.

Figure 4.2: Joint angle definitions. Angle a is defined by a sensor placed on the shoulder, elbow, and end-effector (hand). Angle b is the difference between the hand and the target.

into discrete bins when determining finite observations.

4.1.1.4 Model Construction

A total of 12 HMM pointing models were constructed. Across participants, a unique model was generated for each pointing distance with each of 3 binning sizes (1 degree, 3 degree and 5 degrees) applied to the joint angle data to determine the number of observations. Within each of the aforementioned configurations, 4 state distributions were examined (5, 10, 20, and 30 hidden states). In this examination, all models began with a random initial parameter estimate.

4.1.1.5 Model Assessment

All models were initially tested using an unconstrained MonteCarlo analysis (see Section 3.3.4). Each model was sampled 100 times using this method (observation sequence length 50). This multiple sampling enabled robust central tendency statistics to be calculated for the correspondence between Monte Carlo and recored sequences. A constraint was placed during sampling such that the first observation matched a potential

start position (± 5 degrees or 2 observation values). This constitutes a direct measure of the likelihood that the observed sequence matched those predicted from the measured observation sequences. While this measure is historically considered a central measure for evaluating HMM performance it is unable to account for trajectory similarity, instead merely producing a binary "match-nomatch" score (not unlike a χ^2 or binomial statistic).

In addition to the Hamming distance measure, a Pearson product moment correlation was calculated between the measured and predicted observation sequences. Correlations were estimated for multiple lag values (± 3 samples). The maximum r value was used for further analysis as it represented a more accurate estimate of model-data correspondence.

Summary statistics were calculated for each model (μ, σ), providing a direct estimate of how well the model can encode the training data. Parametric comparisons were carried out for each trial using the z-scores based on the estimated means and variances. For all comparisons involving correlation values, a $r \rightarrow z$ transform was calculated prior to evaluation. The hypothesized "random" score (no prior knowledge) of 1 / *#observations* was used for the Hamming measure and $r = 0$ for correlation measures. A non-parametric comparison (χ^2) was also made for the frequency of significant trials. Chi-square statistics were based on a yes/no grouping from the trial parametric value. Finally, the average Hamming and correlation value were assessed, again using the z-distribution. [‡]

4.1.2 Results

As demonstrated in Figures 4.3 and 4.4 as well as in the results summary (Table 4.1), the hidden Markov models tested successfully demonstrated many of the features required

[‡]In all cases the parametric and non-parametric assessment of model performance were in agreement (parametric value was more conservative), thus only the maximum p value is reported.

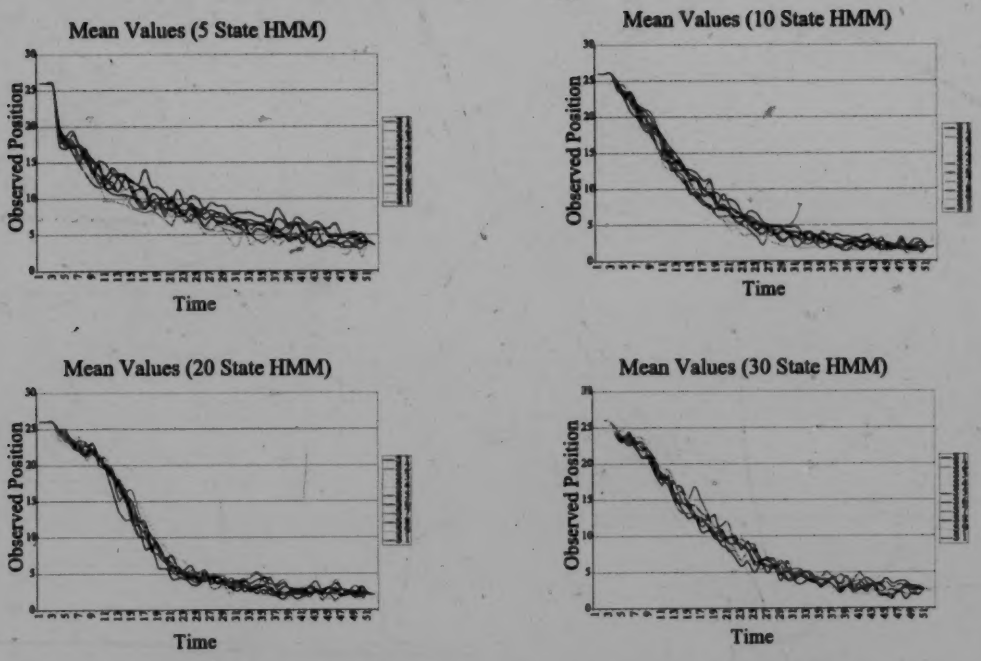

Figure 4.3: Unconstrained MonteCarlo simulation of behavioural data for a 78 degree elbow rotation (3 degree binning size) sampled across all participants. Reported lines represent the average value across 5 sub-sampling runs of 50 Monte-Carlo iteractions.

of any model of aiming. Empirical comparisons of the correlations demonstrated that observed r values were significant ($p < 0.05$) in all conditions except target-two, 10 and 20 state, models. The Hamming scores yielded similar yet more conservative results; here there was a distinct advantage for larger binning values. Five degree binning yielded significant Hamming scores in all cases ($p < 0.05$) and conversely, one degree binning demonstrated no scores reliably different from chance values. More interestingly however, behaviourally 'typical' profiles were generated in many cases; predicted observation sequences subjectively demonstrated features reflective of human performance for both 78^o and 54^o movements (Figure 4.3 and 4.4).[§] Trajectory paths showed path variability while retaining endpoint accuracy (i.e., all predicted trajectories

[§]Due to the variation in production, both medians and means we examined to allow a perspective with regard to the envelope of trajectory likelihood.

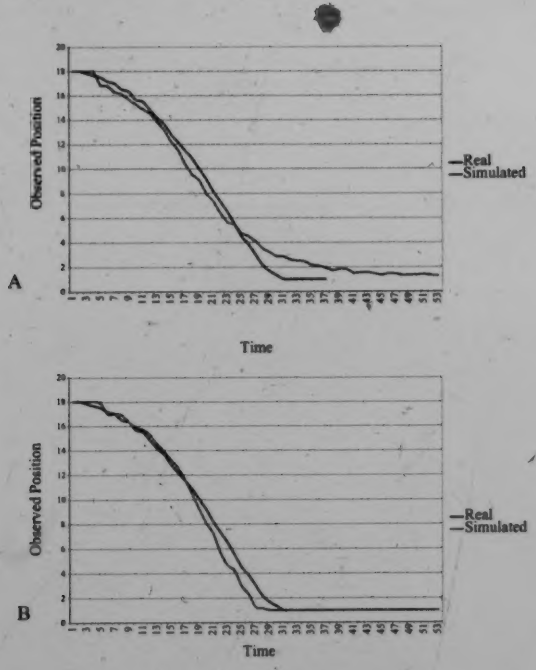

Figure 4.4: Average HMM predicted observation sequences compared to average recorded trials for a target distance of 54 degrees.(a) mean value comparison (b) median value comparison.

attained the target). Similarly, in cases of 20 states (3 or 5^o bins) a visual estimate of velocity indicates profile symmetry; success was highly dependent on the selection of appropriate state numbers and observation clustering.

Although many state-observation combinations within the HMM models presented demonstrated reasonable performance *on average*, individual model are somewhat less appealing. Specifically, predicted observation sequences show 'state perturbations' – unexpected or unreasonable state transitions. Although these occurrences contribute to some of the variability expression, the deviations they suggest are inconsistent with a goal-directed action such as is being discussed here. A further discussion of this point and likely solutions and ramifications of it will be further discussed in Chapter 6. In brief however, these features may be remedied by the inclusion of additional, coupled, HMMs.

4.2 Discussion

HMM performance was promising with regard to its descriptive and predictive application to movement. Results demonstrate the ability of the simple structure of HMMs to represent a dynamical system for pointing. Typical movement trajectories and associated features were represented by a relatively simple combination of transition matrices and observation probabilities. Moreover, no "homuncular node" was required to program the progression of the movement. However, the approach can be criticized as to its relevance toward increased understanding of human motor control. How does the success of HMM methods impact current control hypotheses? How might an HMM be instantiated within the human nervous system and/or constructs hypothesized to dwell therein?

A good starting point for discussing HMM impact may be made by revisiting the original thesis of Berkinblit et al (Berkinblit et al. 1986). In their formulation, the control of multi-joint movements was posited to be produced by independently acting joint automata. Coordination was accomplished by an output evaluation at each unit. While this explanation holds much implicit value, Berkinblit and colleagues did no experimental verification of the control strategy. Due to the formal similarity of HMMs and the automata described by these authors, the current experiment can be viewed as an indirect verification. Additionally, the state control demonstrated by the HMM is performed in a truly stochastic manner instead of relying on inconsistencies of the initial conditions for the differential equations to account for trajectory variability – as is done by many other dynamic accounts of control.

Similar to the position presented by Berkinblit, the symbolic structure of the HMM transition and observation matrices enables a generalization to equilibrium point hypotheses for action. Specifically, the end-state of the joint-muscle system can be captured as an *absorbing* state; thus all progressions of the system will converge on that end-state. However, while the termination may be determined, the evolution of the sys-

tem is free to vary by the stochastic transitions. The end-state can also be expressed by any number of observed features. Thus, while there was no explicit intent to parallel biological systems, the hidden control states being modeled can reasonably be mapped to the EP invariant characteristics.

In addition to Berkinblit however, HMMs can loosely be viewed as a unification between self-organizing viewpoints and the traditional information processing standpoint. This joining is possible by having a dynamical 'packet' (HMM) form the motor program being selected and prepared by the central processor (see Chapter 1). In this way, the 'executive' selects the task type while the representative nature of the HMM reduces the computational load created by trajectory selection and formation.Although it is not reasonable to consider the motor control system performing the Baum Welch update procedure at task initiation, others have successfully argued that the cognitive system performs similar functions, such as backpropagation and gradient descent (e.g., Kawato 1996, Uno et al. 1989) prior to action and perhaps during learning.

In concert with this 'motor programming' explanation of HMM representation comes a possible explanation for the focal weakness of the approach: the occurrence of maladaptive state/observation selections. These events effectively predict that a movement would deviate substantially from the desired/optimal trajectory, perhaps even instantaneously returning to the start position. In fact however, the existence of inappropriate or unattainable commands within the motor system is a known phenomenon. For example, a loaded muscular system is already known to demonstrate low-pass filtering properties, effectively averaging the high-frequency inputs from the α-motor neuron into a manageable command. If one considers a similar process to occur when a series of commands occurs from a given HMM-program, the output is well behaved, but retains flexibility and simplicity.

In summary, HMMs appear functional in the task of movement representation. Although some limitations are apparent, none seem insurmountable within the basic

structure of the modeling method. Similarly, the success of the HMM structure provides some experimental support for local control processes. A demonstration of joint dynamics being accounted for at a joint level is consistent with more general joint control models such as those of the mass-spring proponents.

4.3 Toward a More Generalizeable Framework

The general principle of using HMMs for movement applications seems sound; however, implementing finite-state process for more than a trivial task is awkward. For example, if we were to continue with the premise of Berkenblit, we could, for a multi-joint system, define a set of positions for each joint and combine them into an n-tuple of values representing the system state. There are two obvious criticisms of this generalization. First, as the system grows in size (i.e., more joints) the size of the descriptive vector grows factorially, losing any biological or logical meaning in the process. Second, although the joint level automata explanation was based on independent processors being coupled by their combined output, any independence and/or coupling is lost by the vector concatenation between states.

The logical solution to this state ambiguitity is to construct a coupling property between joint controllers. The generalized dynamical Bayesian network (gDBN), will be examined for its effectiveness at representing just this coupling behaviour. In a three joint reaching limb, this approach to control can be considered by representing each joint as an independent HMM. The interconnections are similarly modeled as Markov chains and their influence mediated by a hypothesis guided weighting factor. From a biological standpoint, the inherent attractiveness associated with the ANN architectures has once again been included. Hypothetically, each HMM might be envisioned as a peripheral controller (perhaps in the spinal cord) with the interconnections being the ensemble of activation at the motor neuron pool level or perhaps as a feedforward/feedback coordinating connection within the cerebellum.

	1 Degree Segmentation		3 Degree Segmentation		5 Degree Segmentation	
	Mean	Std Deviation	Mean	Std Deviation	Mean	Std Deviation
Target 1						
10 States	0.7954688	0.092156488	0.845465182	0.095278933	0.885469462	0.119876543
20 States	0.7690564	0.114765431	0.858569941	0.073443867	0.862262608	0.080347648
30 States	0.81205085	0.074342127	0.878183891	0.113612778	0.858135152	0.083747883
Target 2						
10 States	0.48389801	0.282763191	0.805919	0.137901053	0.772837244	0.122766554
20 States	0.44848191	0.307693829	0.573922297	0.210540025	0.881527804	0.076874391
30 States	0.71649795	0.124947203	0.918165125	0.053548537	0.914952847	0.057140713
Target 3						
10 States	0.60172382	0.244086538	0.603851763	0.243741777	0.886976299	0.07460326
20 States	0.64077467	0.188283884	0.912630064	0.059847697	0.920538612	0.055266429
30 States	0.74893354	0.160542912	0.916450465	0.050988815	0.923219591	0.042431596
Target 4						
10 States	0.64002946	0.218416907	0.871385916	0.080526884	0.864178393	0.078856745
20 States	0.62251208	0.176454538	0.915839796	0.056462252	0.897466224	0.074506236
30 States	0.8791048	0.062284346	0.910036832	0.05829504	0.889853891	0.079729096

Table 4.1: Experiment 1 simulation summary. Results represent the mean and standard deviation Pearson correlations based on a model trained across participants (n = 100, sequence length = 50). Comparison trajectories were randomly selected from participant performance.

Chapter 5

Methods for Increasing HMM Utility

In this chapter we explore how to model complex actions in terms of a generalized Dynamical Bayesian Network, implementing multiply connected nodes (HMMs) capable of encoding both the initial conditions (e.g., position or intention) and emerging behaviours. Such models allow for mixing of hypothesis-based and data-driven parameter searches as well as complex feedback/feedforward relations via connection weights and transitions.

Similar to the presentation of the basic HMM, some discussion of the gDBN will be phrased in reference to a later presented experimental test. Specifically, recall the coupled oscillator models of action discussed in Chapter 2. In these representations the motor control system is regarded as a dynamical system which self-organizes in a fashion analogous to a system of oscillators, describable by relatively compact equations of motion. From the standpoint of modeling this behaviour with a gDBN, one can reasonably hypothesize that two noisy-oscillator HMMs might interact mutually but asymmetrically. The transitions displayed by the human behaviour would be repre-

sented by probabilistic transitions across the connected state dynamics.

5.1 Bayesian Networks

In general, a Bayesian network (BN) may be described as a graphical representation for the (conditional) interdependencies between a set of random variables (Ghahramani 2000). A common application of such a structure is to infer system properties based upon some observed evidence. The process of belief propagation through the BN may be considered as 'evidence gathering' toward the acceptance of a hypothesis regarding the system condition.

For example, we may wish to estimate the likelihood of our car starting. Prior to examining this question, a network must be constructed regarding the interconnections and general relationships between elements of the system (Figure 5.1); e.g., the likelihood of the gas gauge reading "full" is dependent on both the amount of fuel and the charge in the battery (to relay this information). The network structure may therefore simply be a condensing of knowledge or it may suggest some hypothesized formation to be examined. By taking measurements from the battery, the alternator, the pistons (compression), and the gas tank, we may gain or lose confidence in our original proposal (that the car will start). Assuming that observation has determined the car battery to be 5 years old, the goal of belief propagation within a BN is to update the marginal probabilities of the entire network given the new information (e.g., what is the new likelihood of *Engine Start* given an old battery – see http://www.research.microsoft.com/research/dtg/msbn/ for a downloadable demonstration tool).

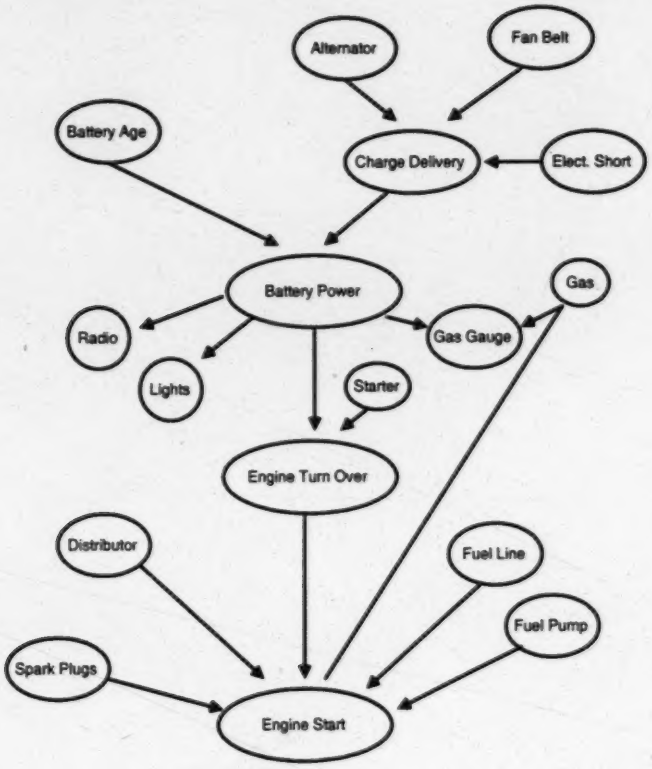

Figure 5.1: Traditional Bayesian network depicting the inter-relations between elements of a car engine.

5.2 Dynamical Bayesian Networks

While BNs are an excellent tool for system modeling and querying, the formulation discussed has no explicit mechanism for capturing system dynamics; a subclass of BNs, the dynamical Bayesian network (DBN), enables such representation. While hidden Markov models, as discussed previously, actually fall into this subclass of BNs, they represent a very simplified representation (Ghahramani 2000). Hidden Markov models summarize time-series with the assumption of first-order causality; in a BN this has conditional arcs flowing forward in time. However, DBN structure may be considered in a more general framework than HMMs. Specifically a dynamical Bayesian network

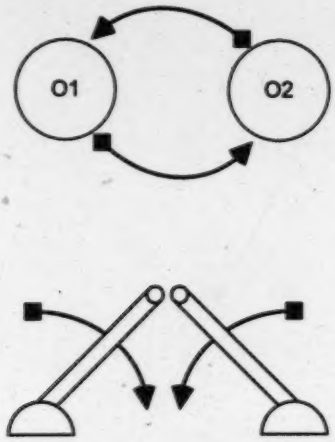

Figure 5.2: Graphical representation of a coupled oscillator gDBN. The two finite-state hidden Markov models are analogous to the oscillators of models such as the HKB-model.

may be constructed by representing the known relations within the modeled system as elements of a 'static' Bayesian network. An HMM is then *inserted* within each node to represent the evolution of that element; Markov chains (A-matrices) are used to represent the dynamics of the inter-relation between nodes. Reflecting on the coupled oscillator example (see Figure 5.2), each component may be expected to behave in a semi-deterministic fashion, usually advancing to the next sequential state. In order to capture interactions, an additional HMM is constructed to represent the dynamics and 'shape' of the interactions between nodes; these inter-node matrices are weighted to represent their strength. The interaction matrices may be considered analogous to the differences in eigen-frequencies hypothesized to underly handedness effects in transition behaviours; the weightings perhaps reflect intention or some similar mitigating property.

5.2.1 gDBN Structure

In general, gDBNs may be considered in a manner similar to HMMs. Just as in the HMM, the model ($\mathbf{\Lambda}$) is parameterized by a transition matrix ($\mathbf{A}$), an observation matrix ($\mathbf{B}$), and the initial likelihood of each state (for all nodes, $\mathbf{\Pi}$). In addition to the

HMM features, a gDBN contains a weighting term ($\mathbf{\Omega}$). The omega term applies a scaling factor to the connecting HMMs, adding a degree of flexibility beneficial in modeling complex system. This plasticity is expressed by the connecting HMM determining the relative likelihood of transitions with a connection and the omega term providing an directed weighting of this influence. However, within a gDBN these components represent compound interactions between smaller HMMs, leading to the requirement of additions to the optimization and analysis algorithms.

Given these additions and extensions, a gDBN may be represented:

$$\mathbf{\Lambda} = \{\mathbf{\Pi}, \mathbf{A}, \mathbf{B}, \mathbf{\Omega}\} \tag{5.1}$$

with

$$\mathbf{\Pi} = \{\pi_u^i = p(S_u^i); i = 1, .., N; u = 1, .., N_i\} \tag{5.2}$$

where N corresponds to the number of HMMs (nodes) in the network, N_i to the number of states in the i'th HMM (Equation 5.2).

5.2.1.1 Generalized A Matrix

The generalized state transition matrix, $\mathbf{A}$, is defined by

$$\mathbf{A} \equiv a_{uv}^{ij} = p(S_v^j(t+1)/S_u^i(t)) \tag{5.3}$$

where i, j correspond to a pair of HMMs and u, v to their states ($u \equiv S_u, v \equiv S_v$). When $i = j$ the state transitions are *within* a given HMM (node). However, when $i \neq j$, the state transitions apply *between* a pair of HMM states(see Figure 5.3 for a 3 node example).* The *shape* of $\mathbf{A}$ encodes knowledge based models by (dis)allowing state transitions between specific HMMs and specific states within and between HMM's.

*In all cases a single unit lag model (t to $t+1$) has been implemented.

Thus, **A** is termed the "causal model matrix" for the generalized DBN.

As is demonstrated in Figure 5.3, the basic configuration of the causal model matrix if fundamentally the same as an A matrix from the one-dimensional HMM. More specifically, each HMM *conditions* the behaviour of the connected HMMs based on the off-diagonal sections – analogous the the connecting arrows in the graphical view (5.3a). From the movement perspective outlined earlier in this chapter (i.e., coupled oscillators), imagine two of nodes to represent each effector and perhaps the third is a metronome timing the movement. From this example one could draw a similar graph: perhaps allowing the limbs to influence each other and to be influence by the metronome – no influence arrows would be drawn to the metronome node. However, while the causal matrix accounts for the expression of interactions between nodes, it does not reflect tenacity of these connection (Ω) or the expression (emergent behaviour) or each node **B**.

5.2.1.2 Generalized Observations (B)

The **B** component of the gDBN serves a similar purpose to the B matrix in a HMM, dictating the most likely observed behaviour for each underlying control state. Thus **B** is a matrix defined by

$$\mathbf{B} \equiv \mathbf{b}_u(o_k^i) = p(o_k^i | S_u^i). \tag{5.4}$$

To be precise, for any state, i, of HMM u (S_u^i), the probability of an observation, k, (o_k^i) is supplied (i.e., given that the oscillator is in state 5, what is the probability distribution for 'finger' positions).

5.2.1.3 Weighting Matrix

The final aspect of the DBN, $\mathbf{\Omega}$, has no analogous HMM term; instead consider a parallel to neural networks. Similar to an ANN structure, the contribution of each

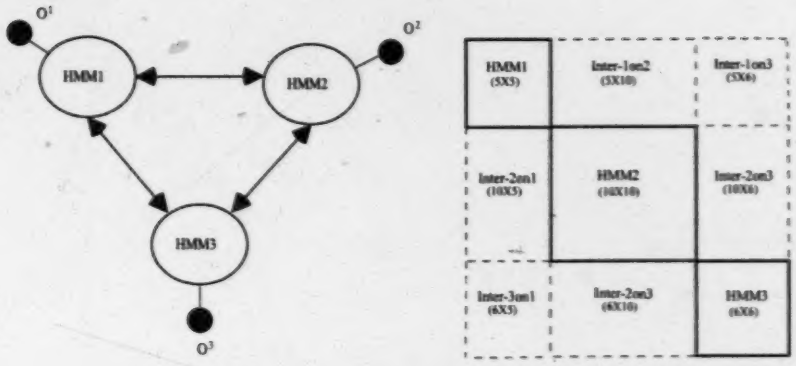

Figure 5.3: Left: Graphical representation of a 3 node gDBN. All nodes are fully connected and may be represented by the causal (Super) A matrix (right). Filled circles represent independently determined state-observations for each HMM node – due to **B**. Right: Causal matrix where the diagonal accounts for within-HMM transitions while the body accounts for interactions between HMM states.

node i to the transition of another node j may be weighted, modulating its systemic influence on the rest of the system. Thus ω^{ij} adjusts the influence of the $a^{ij}th$ element within **A**; if $i = j$ (diagonal) the weighting applies the influence of the HMM on *itself.*[†] Although there is no requirement for $\mathbf{\Omega}$ values to be contained within any range, they are expressed as ratio values. Again returning to movement, this parameter would scale the magnitude of connection between coordinated limbs, perhaps even behaving in a fashion analogous to the b/a ratio from the HKB-model. As yet, there is no update algorithm for optimizing $\mathbf{\Omega}$ beyond the hypothesized initial values. In summary then, a gDBN is simply a series of nodal HMMs linked together by weighted Markov chains. At execution, the state selection (q_{t+1}) of each node is modulated by the current state of all other nodes (based on hypothesized connections). At each time epoch, the expression of each nodal HMM is observed based on sampling of the **B** matrix (see Figure 5.4).

[†]While this term is free to vary, in all of the present examinations this term remains at "1.0" or 100%, reflecting "normal" HMM inter-HMM behaviour. If this term were set to zero it would reflect a virtual observation node; the HMM's transition would have no bearing on its own behaviour but may influence the expression of other nodes.

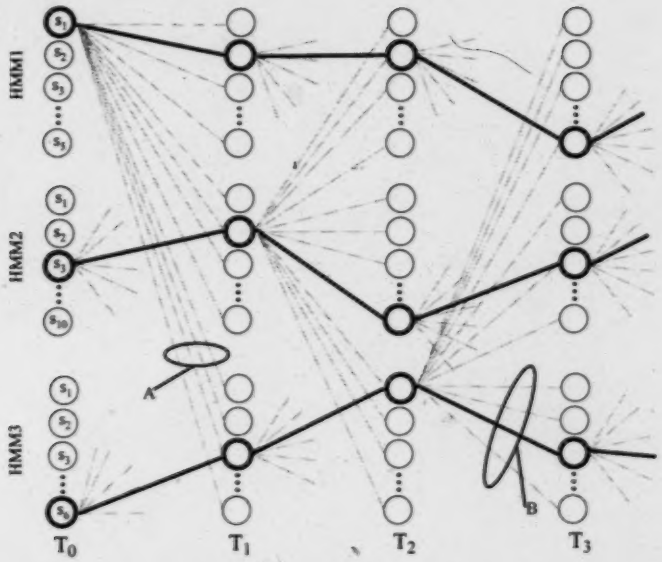

Figure 5.4: Model representation of the 3 node gDBN described in Figure 5.3. The dynamics of the gDBN progress from the selected state for each model(at each T), with those states contributing to the potential transitions for each of the other HMMs at $T+1$. Connection 'weights' are sampled to determine each HMMs individual path (bold lines). This represents a modified mass-action progression similar to the n-Heads formulation (Brand 1997). A: Inter-HMM connections B: Intra-HMM connections.

5.2.2 Estimate Update and Prediction

5.2.2.1 The Generalized Baum Welch Algorithm

As before, the generalized Baum Welch algorithm is a version of Expectation Maximization (EM), updating the initial model estimate based on the observed data. Just as in the one-dimensional case, forward and backward algorithms underly the function of the update procedure. Thus, as an extension of the single HMM case (Rabiner & Juang 1993), the generalized forward operator may be defined:

$$\alpha_1(S_v^i) \equiv \alpha_1^i(v) = p(S_v^i)p(o^i(1)|S_v^i) = \pi_v^i b_v(o^i(1)) \tag{5.5}$$

$$\alpha_{t+1}^{i}(v) = \left[\sum_{i=j,u} \omega^{ii} \alpha_t^i(u) a_{uv}^{ii} + \sum_{j \neq i,y} \omega^{ji} \alpha_t^j(y) a_{yv}^{ji} \right] b_v(o^i(t+1)) \tag{5.6}$$

Here, $\alpha_{t+1}^{i}(v)$ corresponds to the likelihood of a state occurring in HMM v at time $(t+1)$. Following initialization (Equation 5.5), the determination of this probability is broken into a within and a between nodes component. Looking at Equation 5.6, the first component on the right hand side corresponds to the intra-HMM forward operator (identical to single HMM case) while the second refers to the inter-HMM components. Notice that the $\mathbf{\Omega}$ is expressed here as weighting each HMM component in both terms. Accordingly one can model (and restrict) the DBN by blocking specific types of inter- and intra-state transitions (setting the corresponding ω^{ij} to 0).

The generalized backward operator is similarly an extension of the basic HMM operator, divisible into an inter- and intra-HMM component. In matrix form, the generalized backward operator may be expressed as:

$$\vec{\beta}_t = \vec{\beta}_{t+1} \mathbf{A}' \vec{\mathbf{b}}(t) \tag{5.7}$$

For comparison the parallel forward operator:

$$\vec{\alpha}_{t+1} = \vec{\alpha}_t \mathbf{A} \vec{\mathbf{b}}(o(t+1)) \tag{5.8}$$

where $\mathbf{A}'$ corresponds to the transpose of $\mathbf{A}$ and, $\forall ij$, $\beta_T(S_u^i) = 1$.

5.2.2.2 Model Re-estimation

As with the basic Baum Welch (EM) procedure, in the generalized form, a first estimate of the expected state transitions between any two states within a given model is taken. As defined above this incorporates influences of other HMM states within and between the HMMs according to the gDBN causal model. The net result, however, has the same

format as the single HMM case, where, at a given time point, we have:

$$\Phi_t(i,j;u,v) = \frac{\alpha_t^i(u)\mathbf{A}_{uv}^{ij} b_v(o^j(t+1))\beta_{t+1}^j(v)}{\sum_{i,j,u,v}\alpha_t^i(u)\mathbf{A}_{uv}^{ij} b_v(o^j(t+1))\beta_{t+1}^j(v)} \tag{5.9}$$

This amounts to the joint likelihood of a state at time t co-occurring with one at $t+1$, given an observation at $t+1$, for all states and for all time epochs. So, by summing over all states v, for HMMs j we obtain:

$$\hat{\psi}_t^i(u) = \sum_{j,v} \Phi_t(i,j;u,v). \tag{5.10}$$

Here, $\hat{\psi}_t^i(u)$ records the probability of being is a given state v at each time. Consequently, by selecting the first time-point, $t = 1$, a new value for Π can be calculated as:

$$\hat{\mathbf{\Pi}} \equiv \{\hat{\pi}_u^i\} \equiv \hat{\psi}_1^i(u) \tag{5.11}$$

Again similar to the one-dimensional case, by integrating over t we can establish a new estimate for the causal $\mathbf{A}$ matrix (Equation 5.12). That is, we calculate the probability of being in a state v following an instance of state u, across all nodes and times.

$$\hat{\mathbf{A}} \equiv \hat{\mathbf{a}}_{uv}^{ij} = \frac{\sum_{t=1}^{T-1}\Phi_t(i,j;u,v)}{\sum_{t=1}^{T-1}\hat{\psi}_t^i(\mathbf{u})} \tag{5.12}$$

Finally, the likelihood of an observation given a state (i.e., $\mathbf{B}$) can be re-estimated by counting the co-occurrences of observations with states:

$$\hat{\mathbf{B}} \equiv \hat{\mathbf{b}}(o_k^i) = \frac{\sum_{i=1,o_k^i(t)\equiv o^i(t)}^{T}\hat{\psi}_t^i(u)}{\sum_{i=1}^{T}\hat{\psi}_t^i(u)} \tag{5.13}$$

The generalized Baum Welch algorithm is then iterated until a terminating condition is reached; this state may be defined by a maximum number of iterations, a minimum step improvement (as with ANNs), or by a Viterbi estimate of the likelihood of the observed data being produced by the estimated model.

The Generalized Baum Welch Procedure The complete Generalized Baum Welch method may be described as follows:

1. Select a causal model for $\mathbf{A}$ by excluding some dependencies (disallow any interaction not considered reasonable).
2. Generate initial estimates of $\mathbf{\Lambda}$ from the training data, hypothesized structure, or initialize a random network.
3. Re-estimate DBN as $\hat{\mathbf{\Lambda}} = \{\hat{\mathbf{\Pi}}, \hat{\mathbf{A}}, \hat{\mathbf{B}}\}$
4. If $\hat{\mathbf{\Lambda}} \simeq \mathbf{\Lambda}$ STOP
5. Set $\mathbf{\Lambda} = \hat{\mathbf{\Lambda}}$ and GoTo 2.

5.2.3 Assessing gDBN Performance

5.2.3.1 Prediction: The generalized Viterbi algorithm (gViterbi).

As in the case of hidden Markov models, the extension of the Viterbi algorithm for gDBNs predicts the most likely set of state sequences for each node in the system. Given an observation for each node in the gDBN, gViterbi accounts for both intra- and inter-nodal connections, estimating a state vector for the gDBN at each epoch . Like single HMMs this algorithm is based upon computing the most likely trajectories over the "state trellis" using a forward estimate of likelihood of each state at each time (given the observations). This set of state arrays, one for each node, is then searched to determine the optimal solution.

Algorithm

For the most part, the gViterbi algorithm proceeds similar to the Viterbi procedure. Specifically, at initialization δ (the likelihood of a given state based on the previous

epoch) for each state of each node is determined from the π (the overall likelihood of a state) and $b_u(o^i(1))$ (the probability of the observation given that state).

Initialization

$$\text{For } 1 \leq i \leq N,\ 1 \leq u \leq N_i$$
$$\delta_1(S_u^i) \equiv \delta_1^i(u) = \pi_u b_u(o^i(1))$$
$$\Phi_1^i(u) = 0$$

Recursion

The combined likelihood of each state, given all antecedents at each node is estimated and held in an array (δ); the maximum estimate for each is selected and stored (Φ). These arrays, analogous to the δ and ψ vectors previously, are determined by the likelihood of transitions from each of the previous states $(t-1)$. As was presented in the discussion of the generalized Baum Welch procedure, all connections are conditioned by an ω^{ij} factor – as determined by hypothesis or other prior knowledge.

So, for $2 \leq t \leq T$, $1 \leq j \leq N$, $1 \leq v \leq N_j$, recurse as:

$$\delta_t^i(v) = max_u[\delta_{t-1}^i(u)\omega_{ii}a_{uv}^{ii} + \sum_{j \neq i,y} \delta_{t-1}^j(y)\omega_{ji}a_{yv}^{ji}]b_v(o^i(t)) \tag{5.14}$$

$$\Phi_t^i(v) = argmax_u[\delta_{t-1}^i(u)\omega_{ii}a_{uv}^{ii} + \sum_{j \neq i,y} \delta_{t-1}^j(y)\omega_{ji}a_{yv}^{ji}]b_v(o^i(t)) \tag{5.15}$$

The resulting "state trellis" encodes the maximum likelihood of being in each state of each HMM at a given t.

Back Tracking

Once again, following the construction of a an array of the probability of each state for all epochs (given an observation) the backtracking routine selects the most likely state for each time based on state-transition probabilities (i.e., **A**). The procedure continues from time T to t_1. For all HMMs, i:

$$u^i(t-1) = argmax_v \left[\Phi_t^i(v)\right]; T \geq t \geq 2 \tag{5.16}$$

It is important to note that the trellis search is of size TCN for sequences of length T and C HMMs each having N states, thus reducing the factorial search explosion occurring in exhaustive search. This benefit is due to the "mass action" model explored here[‡]; combinatorial vectors of states between the HMMs were not examined (i.e., estimating δ for each of the possible vector combinations of states across nodes). As will be discussed in Section 5.3 this is not the only possible implementation of this framework, however, it does limit the number of required simplifying assumptions, thus allowing more functionality.

While the basic Viterbi algorithm was discounted earlier as a controller[§] it does provide an excellent opportunity for initial examinations of complex systems. Specifically, the Viterbi-derived posterior maximum likelihood probability (traditionally reported in *log* units) can be used for trajectory recognition; this value reflects the probability of the observed data sequence arising from the selected model.

As will be discussed further in Chapter 7, the gViterbi approach can be used for establishing estimates of noise and coupling parameters in a coupled oscillator model of reaching. Briefly, for now, gViterbi can evaluate the degree of fit of a model to a data source; in a coupled oscillator scenario, this approach would allow the evaluation of the suitablility of the pre-transition coupling versus post-transition value for coordination behaviour that may be ambiguous (this is of course presuming a 'shift' in the underlying control parameter).

5.2.3.2 Monte Carlo Sampling of gDBNs

Similar to the procedure discussed earlier for the HMM case, for each node we can generate a predicted observation sequence by randomly selecting observations according

[‡]Mass action is a term that simply reflects the combined influence of all terms, as in ANNs.

[§]Viterbi was set aside due to its requirement of a complete observation list prior to initiation. This is not a reasonable feature of a realtime controller (e.g., discrete aiming/manipulation). However, it can be used in an "incremental" mode for control applications – see Chapter 8

to the estimated DBN model probabilities. This amounts to an unconstrained Monte Carlo sampling of the model distributions such that:

$$\textit{choose } \vec{x} \propto p(\vec{x}) \tag{5.17}$$

$$\textit{choose } \vec{o} \propto p(\vec{o}|\vec{x}) \tag{5.18}$$

However, similar to the single HMM case, a constrained Monte Carlo method may just as easily be implemented: instead one simply selects from observations given the optimal (gViterbi) state sequence. Regardless of constraint, for the case of the gDBN, this procedure can become quite complex. Each HMM component, random-variables, and associated coupling variables, can be Monte Carlo sampled using various hypothesis driven constraints. These may include the most direct interpretation of Monte Carlo sampling: sampling each variable and then using this variable to condition sampling all other dependent variables, to eventually select the observation. However, without making extensive simplifying assumptions, sampling such a set of multivariate densities proves a difficult task due to the increasing number of alternatives. For example, in an model consisting of 5-HMMs with state segmentations of 5, 10, 6, 4, and 9; The number of possible state vector configurations is $10,800$.

Consistent with the Viterbi and Baum Welch algorithms, a Monte Carlo sampling mode was developed which is consistent with our 'mass action' formulation for gDBNs. That is, each model is "fully connected" with all other nodes, provided weighted influence at each epoch. Thus for each HMM, i, we have:

$$x^i(t+1) \propto \left[\sum_{i=j,u} \omega^{ii} \alpha_t^i(u) a_{uv}^{ii} + \sum_{j \neq i,y} \omega^{ji} \alpha_t^j(y) a_{yv}^{ji} \right] \tag{5.19}$$

Again consistent with the previous approach to model sampling, this approach first randomly selects a new state for each node (i) based on the probability of each state

(u) given all other models (Equation 5.19). Given a state for each HMM, now generate an observation by sampling the **B** matrix:

$$o(t) \propto p(o|s(t)) \tag{5.20}$$

5.3 Other Approaches to gDBNs

Although the previous sections imply a single algorithm for gDBN estimation, other approaches have been suggested (Brand, Oliver & Pentland 1997, Brand 1983, Caelli, Binsted & Frankland in submission, Ghahramani & Jordan 1997, Ghahramani 2000). These methods take a range of approaches. However, researchers, for the most part, attempt to reduce the update problem by utilizing 'cliques' of states within the connected HMMs. This approach can be considered analogous to using the δ terms from the Viterbi algorithm for model update. For the purpose of reflection and comparison, two alternate approaches to gDBN update are briefly presented; each bears a marked similarity to the approach described here.

5.3.1 N-Heads Approach

$$MassAction \longrightarrow \alpha^i_{t+1}(u)\left[\omega_{ii}\sum_{u;i=j}\alpha^i_t(u)a^{ii}_{uv} + \sum_{w;i\neq j}\omega_{ij}\alpha^j_t(w)a^{ij}_{uv}\right] b_v(o^i(t+1)) \tag{5.21}$$

$$NHeads \longrightarrow \delta^i_{t+1}(u)\left[\omega_{ii}\sum_{u;i=j}\delta^i_t(u)a^{ii}_{uv} + \sum_{w;i\neq j}\omega_{ij}\delta^j_t(w)a^{ij}_{uv}\right] b_v(o^i(t+1)) \tag{5.22}$$

Brand (1997) presents an approach very similar to the current "mass-action" implementation, there are some notable differences; the most prominent difference being the method of probability calculation prior to Baum-Welch recursion. In the N-Heads implementation a δ term (see Section 3.3.3.3) is summed; this amounts to using the

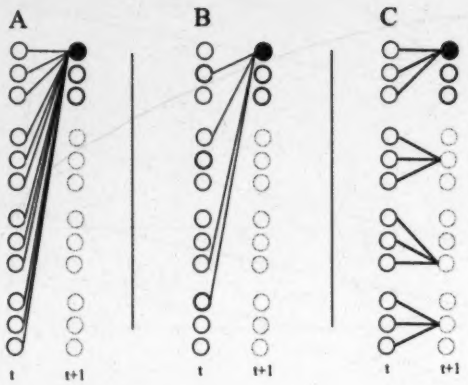

Figure 5.5: Schematic of the (A) Mass action (B) N-Heads and (C) Residual approaches for generalizing the Baum Welch algorithm.

'most likely' path only during EM search. By making such an assumption the N-Heads approach reduces the number of sequence alternatives, selecting state transitions based on state "cliques" (see Figure 5.5).

Conversely, in the mass-action implementation used in this investigation, no path assumptions are made during the Baum Welch algorithm[¶]. This difference has several implications for the update procedure. First, while the N-Heads procedure does limit the search to the 'most likely' contributors it is unsuccessful in similarly reducing the computational load. Second, while the performance advantages for the mass-action versus N-Heads approaches are disputable, and likely dependent on the application or data, there can be little argument of the clarity of generalization provided in the current method. The elements contributing variance in the mass-action formulation of the EM estimate are separable, allowing further additions or restrictions to be easily implemented (e.g., an additive model of node interaction).

[¶]A similar assumption to the N-Heads model is however utilized during MonteCarlo sampling.

5.3.2 Residual Approach

In parallel to the formulation presented in this chapter an alternative "correlational" or "residual" algorithm was investigated for capturing the degree to which the intra- and inter- state dependencies are expressed, in parallel or sequentially (Caelli et al. in submission). The basis of this approach was to estimate the HMM behaviour of each node independently, utilizing traditional HMM update procedures. Subsequently, HMMs were trained to represent the interaction between nodes, with the observation of each intra-HMM being the state of the inter node (Figure 5.5). While not a as complete a representation as the previously discussed models, this approach has benefits. In particular, given an unknown system structure (i.e., no prior knowledge), the relations between node expressions can be "discovered" with a known update procedure. Alternately, in circumstances of reduced data precision, such a methodology could be used for data supplementing (estimating missing data given neighbouring node states). However, while this approach does simplify model estimation by reducing the "space" to be searched via EM, it also has the drawback of being unable to simultaneously search inter- and intra- connections to find the optimal balance.

5.3.3 gDBNs for Action

Although a number of algorithms are possible for optimizing a generalized dynamical Bayesian network, it is not the intricacies of the method which make the gDBN attractive for modeling human data. Unlike many other modeling methods, the node based underpinnings of the gDBN allows researchers to "add dynamics" to models previously considered to be static (e.g., "information-processing" approaches to action). For example, recent behavioural examinations of motor tasks suggest a common timing mechanism housed within the cerebellum (e.g., Ivry & Keele 1989). This structure is proposed to interact with the motor system, enabling precise motor synergies. However, behavioural researchers implementing traditional IP approaches, while able to

demonstrate overt effects, are unable to accurately represent the dynamics of the interactions. Similarly, proponents of non-linear modeling techniques such as those proposed by Haken and Kelso (Haken et al. 1985) have difficulty to capture this relation in more than abstract terms. This topic is pursued further in Chapter 6.

Chapter 6

Dynamical Bayesian Networks in Action

As was discussed formally in Chapter 5, generalized dynamical Bayesian networks present a unique tool for modeling complex dynamical systems. In this and the following chapter, several examples of applications and implications of this approach for modeling human action are presented. First, an extension of the previously discussed one-dimensional pointing model is described. Later in the chapter a theoretical extension of the 2-HMM pointing model is applied to a coupled oscillator perspective on continuous (phasic) human action using a gDBN foundation.

6.1 Experiment 2: A Time-Indexed Reach

6.1.1 Introduction

Earlier, an experiment was presented which tested the applicability of a single HMM for reproducing one dimensional pointing movements. One of the central criticisms of this

approach to motor programming was the presence of "state perturbations" – instances where the model predicted impossible or impractical state selections. Secondly, the HMM was oversensitive to simple variations in amplitude or temporal scaling; similar to early notions of motor programs (Schmidt 1982), the HMM representations were unable to generalize across movement parameters (e.g., movement duration). Relative scaling of observations and states could have overcome these problems, by encoding temporal and spatial relations for each task; however, such solutions lack relevance to biological function (i.e., the novely problem). The following experiment suggests an alternate solution.

6.1.1.1 A Metronome for Action

The 1989 proposal of Richard Ivry and Steve Keele (Ivry & Keele 1989) for an internal timing mechanism for the motor system provides an interesting solution for the stability and generalizing problems encountered by the earlier HMMs. Specifically, the hypothesis that a central timing mechanism is housed within the cerebellum has received significant support both behaviorally (e.g., tapping; Ivry & Keele 1989, Wing 1973) and through clinical studies (e.g., cerebellar lesions; Ivry, Keele & Diener 1988, Franz, Ivry & Helmuth 1996).

So, reconsider the case of simple pointing. If a central timer does exist, as unique process intimately linked to the controlling system, this process would impact the state dynamics of that system. Moreover, a single HMM would be unable to represent this compound system and its temporal interactions, as one of the underlying assumptions of the HMM method has been compromised. Therefore, perturbations and similar errors in a system lacking a timekeeper are reasonable to expect. Similarly, the presence of a timekeeper would introduce a degree of generality by discriminating tasks based on expected duration (or relative duration; see Generalized Motor Programs in Chapter 2).

6.1.1.2 A gDBN for Reaching

In keeping with the discussion of the potential benefits of an integrated time keeper, a generalized DBN was constructed with such a feature. In particular, a second node (HMM) was added to the generation model featured in Chapter 4. This structure was examined for its ability to reproduce joint angles measured during manual pointing.

Due to the preliminary nature of this examination of gDBNs for representing biological processes, some assumptions were made regarding interconnectivity. Although there are undoubtedly significant cerebellar inter-connections with the planning and controlling regions for this task (e.g., pre-motor cortex) only unidirectinal connections from the cerebellar "clock" to the control mechanism were allowed. This constraint reflects the assumption that the timekeeper is generalized in some fashion and therefore should be independent in function from the processes it times. Also, in keeping with the Markov assumptions underlying gDBN structure, it is assumed that the process of all nodes within the gDBN operate on a single temporal scale. Within the gDBN framework, these features were expressed within the Ω term of the model; the connection term representing controller influence on the timekeeper was 'dissallowed".

6.1.2 Methods

6.1.2.1 Participants

The same four male subjects as presented in Experiment 1 volunteered again to participate in this investigation (mean age = 26 years). All participants were right handed and naïve to the purpose of the experiment.

6.1.2.2 Task and Procedure

Each subject completed a total of 20 elbow-pointing trials with their right arm to each of 4 targets (approximate mean angles of 20, 50, 75 and 110). In all other manners, the task and instructions were identical to Experiment 1 (Chapter 4, Section 4.1.1).

6.1.2.3 Apparatus

The experimental equipment, collection, and data reduction procedures were identical in all manners to the first experiment (Section 4.1.1).

6.1.2.4 Model Construction

Once again, the goal of data modeling was to reproduce the observed angular positions. Unlike Experiment 1 however, only one binning size (3 degrees) and one hidden state configuration (25) were used for constructing the HMM representing the elbow controller. This state structure was chosen as a mid-point for all the to-be-learned actions. The 3-degree observation binning size was deemed most appropriate as a tradeoff between observation continuity and temporal scale.* DBN models were trained on random sampling of trials from all movement eccentricities. Again similar to previous methods, participant data was pooled to attain a model characterizing "normal" function for simple reaching.

The clock component consisted of a 40-state HMM with an equal number of observations deterministically dependent on the corresponding state (i.e., a finite state machine)†. A range of connection strengths for the clock node were examined (coupling, ω, see Chapter 5), each composing a new gDBN model. Specifically, unidirectional weights of

*This assumption should be avoidable in future DBN models as an extension to both the Baum Welch and Monte Carlo routines is possible for a continuous observation distribution.

†This allowed us to encode a simple deterministically advancing clock into the DBN format.

100%, 10%, 1%, and 0.1% relative to the within process were examined (i.e., the action HMM was dependent on the clock but not vice-versa). Given that no observations could be made for the clock variable, a pseudo-deterministic structure was hypothesized for its dynamics[‡]. An HMM with this 'shape' was used to generate an observation sequence for training purposes.

Both the clock and control HMM's were trained on time-matched sequences 40 samples in length; this was sufficient to capture all training movements. The only additional restriction placed on training was to initiate the clock starting state to one reflective of the movement time of the trial being 'observed'.

6.1.2.5 Model Assessment

Model evaluation was undertaken in three modes:

(1) Known Sample - Observed Amplitude and Duration

Initial tests were performed by comparing model output with a known mid-range movement (75 degree movement). As before, all models were tested using an unconstrained Monte Carlo analysis (see Section 3.3.4). As an extension of the single HMM method, sequences for all internal HMMs were concatenated to produce a single sequence reflective of total performance. Each gDBN was sampled 100 times using this method (sequence length 50).

As in Experiment 1, a Pearson product moment correlation (see Experiment 1) was calculated over multiple lags (± 3 intervals) to remove phase offset biases. In addition, an observation dependent distance measure, similar to a discrete RMS error, was also examined. The *mean-state-distance* (MSD) was calculated for each model sequence as

[‡]A 95% structure was assigned. This structure entailed that the process remained in the current state 2.5% of the time, advanced one state 95%, and advanced 2 states 2.5%. As no reasonable alternative hypothesis was available, these states were expressed by their corresponding observation 100% of the time.

would be expected:

$$MSD = \frac{1}{N}\sum_{i=1}^{N}|d_i| \tag{6.1}$$

where $d_i = (obs_1(i) - obs_2(i))$ – that is, the difference between the actual and observed sequence at each epoch. This returns a range of values which reflects the distance between sequences, not just a binary 'yes/no' comparison.

Summary statistics were calculated (μ, σ), capturing the performance and variability of the model. A parametric comparison using the z statistic was also applied; a hypothesized null score of 0.04 $(1/\#observations)$ was used for the Hamming measure, $r = 0$ for correlation measures[§], and 12.5 ($\frac{1}{2}$ observation range) for the MSD measure. A non-parametric comparison (χ^2) was also made based on the frequency of significant trials for each measure. Chi-square statistics were based on a yes/no grouping from the trial parametric value. Finally, the average Hamming and correlation value were assessed, again using the z-distribution.[¶]

(2) Novel Amplitude

Model scaleability was examined by evaluating the efficacy of the model for generating unobserved movement amplitudes. To accomplish this, the initial state of the control ("arm") HMM was set to a 'short' and a 'long' value. The short amplitude roughly corresponded to a 30^o movement and correspondingly the long to a 60^o movement.

(3) Novel Duration

The second aspect of scaleability – time – was examined here. Specifically, the initial state of the clock HMM was set to one of three values. The trained value (hereafter 'accurate') corresponded to normal participant instructions; models were trained on trials where subjects were told to point to the target accurately and rapidly. The model

[§]Pearson product correlations were transformed to a z-distribution prior to analysis.

[¶]Again, similar to the previous examination, in all cases the parametric and non-parametric assessment of model performance were in agreement (parametric value was more conservative), thus only the maximum p value is reported. In all instances, each dependent measure was in agreement as to detecting a significant/non-significant dichotomy.

here was sampled according to normal model start parameters. An unobserved fast and accurate (speed as a priority over accuracy) and a fast condition were also examined. In these hypothetical scenarios the clock HMM was set to either an intermediate or a short duration, simulating an intention to move more quickly.

Model performance under conditions 2 and 3 were evaluated based on trajectory shape and appropriate movement scaling in accordance with the temporal or spatial manipulation of the model.

6.1.3 Discussion and Results

Overall, participants performed the required pointing action with a movement time of 561 ms ($\sigma = 62ms$) for the modeled reaching movements of 75 degrees. Performance on all other amplitudes were of similar order and are displayed in Table 6.1.

As illustrated in Table 6.2, model performance was again very good. The Hamming score, Pearson-r, and MSD values reached levels significantly different from chance ($p < 0.01$) in all conditions except the 100% weighting condition. In this condition, performance was at chance level for all measures. These results are consistent with impressions reached by visual inspection of Figures 6.1 to 6.3. When coupling strength was greater than 10%, the model became untrainable; producing a structure that predicted immediate transition to the termination state (Figure 6.1). However, model performance rapidly improved even with only slight decreases in coupling strength. Further, as coupling continued to weaken, participant performance degraded to the performance demonstrated for an uncoupled controller (Experiment 1). This decrease is reflected in a rapid increase in observation variability as state perturbations re-emerge (Figure 6.3). Of particular importance is the finding that a clock mechanism, in combination with the joint based controller, stabilized performance.

Also of note are the performance differences due to variations in clock connection

	Target 1 (20 deg.)	Target 2 (50 deg.)	Target 3 (75 deg.)	Target 4 (110 deg.)
Movement Time	379.17 ms	483.33 ms	561.11 ms	602.65 ms
(stdev)	28.46 ms	55.78 ms	62.06 ms	87.18 ms

Table 6.1: Summary of participant performance for Experiment 2. Movement times are the grand-mean, averaged across participants and trials while the standard deviation values represent the average variability within a subject, averaged across subjects.

Target		100 % Connection Weight	10 % Connection Weight	1 % Connection Weight	0.1 % Connection Weight
75 degrees **	Pearson - r	0.0761613	0.910922	0.950262333	0.898814092
	Mean State Distance	11.438	3.301	2.814	3.41

** Note: 75 degrees corresponds to a 25 unit movement (3 degree bin size)

Table 6.2: Summary of results for Experiment 2. Model performance is reported for the 75 degree movement only. Both correlation and Mean State Distance measures are the maximum value returned after ± 3 time interval phase shifts.

strength. As is evidenced by the inability of the 100% model to acquire an appropriate trajectory shape, we can now make a reasonable hypothesis regarding the relative weakness of any clock-controller interactions (again, see Figures 6.1 and 6.2). The predicted trajectories would suggest a maximum relative connection strength somewhat less than 10% and a minimum strength of greater than 0.1% (as performance appears to be degrading at this point). This proposal, while based solely on the ability of the simple model presented, provides a testable question to be addressed by further behavioural or neuro-anatomical investigations. In particular, this result may have relevance when considering ablative injuries to the cerebellum (or similarly expressive apraxias). In such situations, behaviour tends to become dismetric – a feature very similar to the earlier examination of a single controller in the absence of a timekeeper. Perhaps also, in conditions of less severe cerebellar intrusions (e.g., clinical cooling) the hypothesis could be posited that instead of the clock itself being disrupted, the connection strength is instead being adjusted (perhaps increased). At a deeper level, these results demonstrate that if a "timer" or "internal-clock" is explicitly involved in the execution of motor tasks, then its influence is quite specific. That is, it plays a necessary "background" role such that if its influence is made to play a more dominant role, performance would become inappropriate (e.g., Figure 6.3).

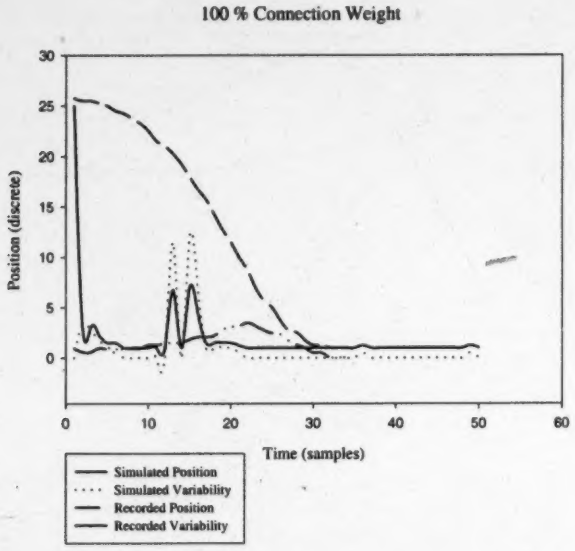

Figure 6.1: gDBN simulation using a coupling strength of 100% between the clock and controller for the training amplitude of 75 degrees (25 units, binning size of 3°.)

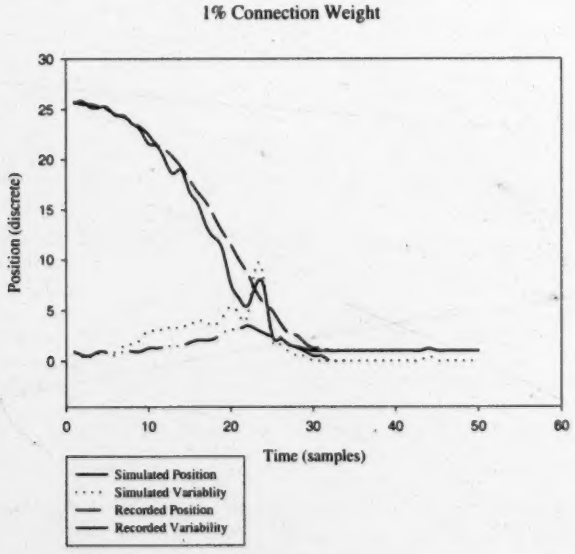

Figure 6.2: gDBN simulation using a coupling strength of 10% between the clock and controller for the training amplitude of 75 degrees (25 units, binning size of 3°).

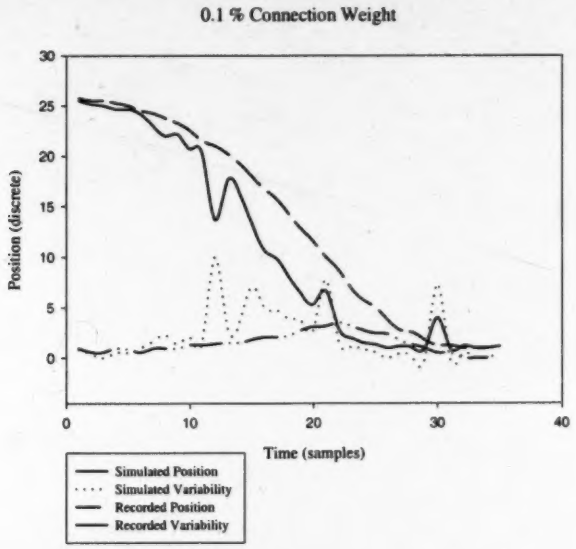

Figure 6.3: gDBN simulation using a coupling strength of 0.1% between the clock and controller for the training amplitude of 75 degrees (25 units, binning size of 3°).

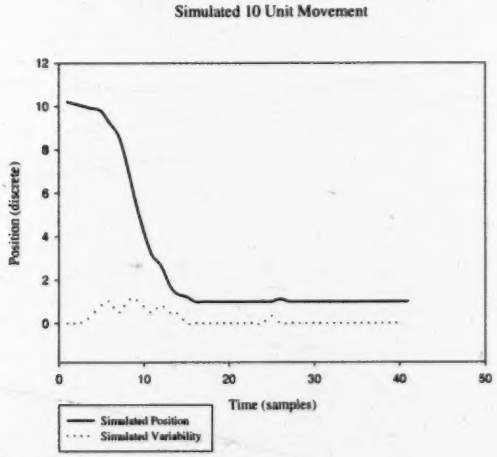

Figure 6.4: gDBN simulation using a coupling strength of 1% between the clock and controller for an "unseen" amplitude of 30°.

6.1.3.1 Model Generalization

In addition to reducing the number of state perturbations, it was also hypothesized that the model would become more generalizeable for both time and amplitude modulations

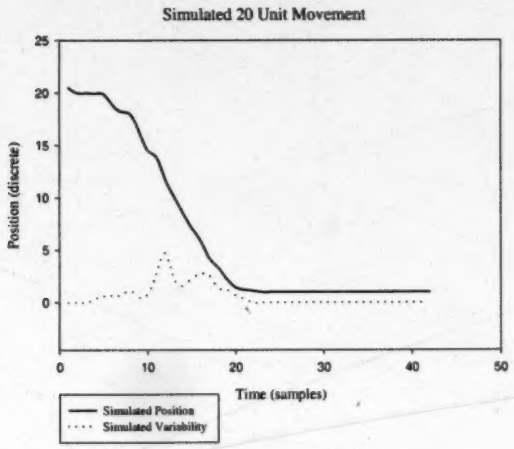

Figure 6.5: gDBN simulation using a coupling strength of 1% between the clock and controller for an "unseen" amplitude of 60°.

with the addition of a 'clock' node. As is apparent in Figures 6.4 and 6.5, simulated pointing trajectories appear reasonable across quite a wide range of amplitude; this variation was accomplished solely by varying the start position of the limb-control (amplitude) intra-HMM. Unfortunately, while mean performance does appear reasonable, with both 10 and 20 unit "movements", demonstrating appropriate trajectory shapes, trace variation is suspect. Specifically, movement duration increased with amplitude as expected while mid-movement variability was unexpectedly high, although not completely unreasonable. Moreover, the short amplitude movement variability remained low throughout trajectory, even with the increased movement 'velocity'. This final result is somewhat at odds with Fitts' Law (Chapter 1), perhaps suggesting a limitation to the simple model structure chosen here.

Modeling performance under the time-variation conditions was also encouraging. As is demonstrated in Figure 6.6, variations in trajectory did result from artificial movement "intent" to move more quickly. Specifically, the 'fast' and 'accurate' conditions are noticeably different, bearing features suggestive of differences in movement variability and perhaps even reaction time. Although not clear, the intermediate 'fast and accurate' condition does appear intermediate to the other conditions. Unfortunately again,

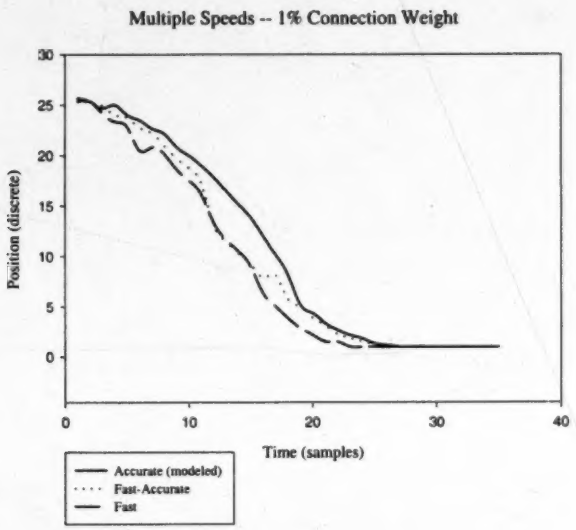

Figure 6.6: gDBN simulation using a coupling strength of 1% between the clock and controller for 3 temporal goals: accurate aiming (training), fast and accurate aiming, and fast as possible aiming.

variability was not ideal; all three conditions demonstrated near identical patterns and magnitudes of variability.

6.1.4 Conclusions

The results reported in this cursory examination of gDBNs for representing single joint movements suggest a significant potential for such methods to be used in predicting and statistically prototyping discrete tasks. Similarly, model behaviour strongly alludes to the stability and flexibility of the predictions possible by combining several simple units – a sentiment that is strongly supported in Berkinblit's 1986 commentary (Berkinblit et al. 1986) – suggesting many new directions of investigation. However, while there may be sympathetic explanations from a single joint standpoint of motor control (e.g., Bizzi et al. 1976, Feldman 1986) and a reasonably easy conjoining with information processing control proponents (Fitts 1954, Meyer et al. 1988, Rosenbaum et al. 1999), a reasonable question that arises is how this approach relates to current dynamical

systems theories for discrete tasks (c.f., Schoner 1990, Zaal et al. 1998). More generally, how does this method generalize to continuous tasks? The reverse of this extension is a central criticism of dynamical pattern perspectives. Can gDBNs account for the demonstrated non-linear behaviour in coordination tasks? – a feature so easily and *simply* accounted for in these theories.

In the next experiment, a hypothetical system of coupled oscillators is framed within a dynamical Bayesian network, examining the efficacy of the gDBN approach for capturing just this type of continuous coordinated activity. More specifically, the goal of the next investigation is to examine ability of the gDBN method for representing continuous tasks, in particular the features of non-linearity (hysteresus, phase transitions, etc.) typical in examinations reported in support of the seminal work of Scott Kelso (e.g., Haken et al. 1985, Kelso 1984).

6.2 Experiment 3: Hypothetical Coupled Oscillators

6.2.1 Introduction

As discussed in Chapter 3, the basis of dynamical pattern perspectives toward movement organization are founded in the ideas of self-organization and complex system behaviour. Researchers endeavor to present control systems for coordinated behaviour that reduce the effective degrees of freedom problem (Bernstein 1967) by coupling joints, muscles, or stimuli into movement synergies or coordinative structures. While the seminal work in this area was produced for simple oscillator behaviours such as finger or wrist oscillations (Haken et al. 1985, Kelso 1984), applications of this approach have been used to examine wide ranges of behaviours (e.g., gait) as well as accounting for abstract contexts such as learning and intention.

In an initial attempt to reproduce the behaviours typically shown for continuous co-

ordination tasks and thus at least superficially replicate dynamical systems results, a model composed of two variably coupled oscillators was constructed. The precise form of coupling, both in strength and determinism, was varied widely to better reflect the possible expression of such a system (see Figure 6.7 and Tables 6.3 to 6.5). This range examination should be considered a close approximation to similar investigations of the Haken-Kelso-Bunz model where coupling and noise parameters were scaled (Byblow et al. 1998, Carson 1995). Unlike the earlier presented modeling experiments with HMM and gDBN structures, in this investigation there is no "ground-truth" data to compare with model predictions to (as we wish to examine behaviours that might not be realistic but may be enlightening). Moreover, owing to the large literature regarding continuous coordinated behaviour there is a reasonable expectation that a model demonstrating features such as phase transitions and hysteresis has some similar properties to a human coordination system. While it is acknowledged that extant coupled oscillator models successfully attend to a wider variety of features than these, transitions and hysteresis are the most salient observable behaviours and are thus appropriate to examine here.

6.2.2 Model Construction

Three hypothetical models were constructed and tested. For all HMMs, 10×10 state transition matrices were used, with a similar 10×10 observation matrix (i.e., 10 observation and 10 states). In order for the underlying state behaviour to be directly observable this observation matrix was deterministically related to the states.[‖] The final common feature for all models was the range of connection strengths and symmetries examined (Table 6.3, Ω). In these 2×2 matrices the diagonal values represent the inter-HMM strengths, these values are always assumed to be 1 as there is currently no legitimate reason to expect otherwise; this does not mean that oscillators are assumed

[‖]While the deterministic relation between state and observation does make the process more transparent, it does limit the degree of nonlinearity possible in model expression. This relation also effectively reduces model behaviour to that of a pair of coupled Markov chains

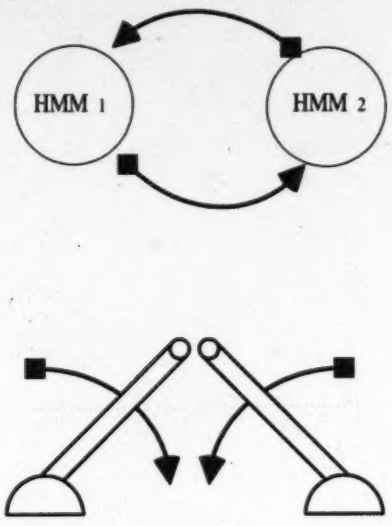

Figure 6.7: Graphical representation of coupled oscillator model for continuous action. Oscillators are taken in this representation to account for cyclic behaviours such as the "finger-wiggle" experiment so popularized by Scott Kelso.

to have identical properties (this difference would be captured in either the A matrices or the inter-HMM structure). The off-diagonal values represent the inter-connections between oscillators/nodes, one place to examine symmetry and strength of coupling.

In a similar fashion to the previous examinations of both HMMs and gDBNs, the models were Monte Carlo sampled to produce a predicted observation sequence for both nodes in the system. This sampling was carried out consistent with the description in section 5.2.3, basing output observations on the inter-HMMs, intra-HMMs, and the assigned connection weights. Due to the restricted, discrete-time, nature of this investigation, relative phase was estimated simply by a temporal ratio relative to each cycle duration. Thus a value of 0.5 corresponds to anti-phase behaviour ($\pm 180^o$) and 0.0 to in-phase (0^o or 360^o).

6.2.2.1 Deterministic Oscillators

The first model was built to examine the coordinative properties of a *perfect* (i.e., noise free) oscillator system (see Table 6.3). As such the transition matrix of each inter-HMM node was set as deterministically advancing. Thus, if these two oscillators were to be controlled independently, regardless of the initial phase relation, they would be expected to remain in that relation.[**††] While the shape of the inter-HMM remained constant for all examinations of this model, the connection weight manipulations were expected to scale the importance of this shape to overall model expression (see Table 6.3).

6.2.2.2 Semi-deterministic Oscillators

This second model was an arbitrary selection of a low-noise system. While many possibilities existed for how such a system might be expressed, the 5% noise level was considered a reasonable hypothesis for a biological signal-noise relation (see Table 6.4). Thus, if these two oscillators were to be controlled independently, regardless of the initial phase relation, they would be expected to undergo random "phase wandering" due to their inherent variability. The inter-HMM A-matrices were similarly established with 5% noise.

6.2.2.3 Noisy Oscillators

The final model considered was a high-noise system. The noise level introduced here, 20% across 3 neighbouring states, was established based on observation of independent oscillator behaviours (see Table 6.5). In these pilot examinations, the 20% configuration

**The in-phase / anti-phase distinction throughout this investigation is of course arbitrary and the models could similarly have been arranged to destabilize the in-phase relation

††Inter-HMM shape was set as pseudo-deterministic here, this was based on the lack of a reasonable hypothesis regarding the deterministic interaction between two deterministic processes.

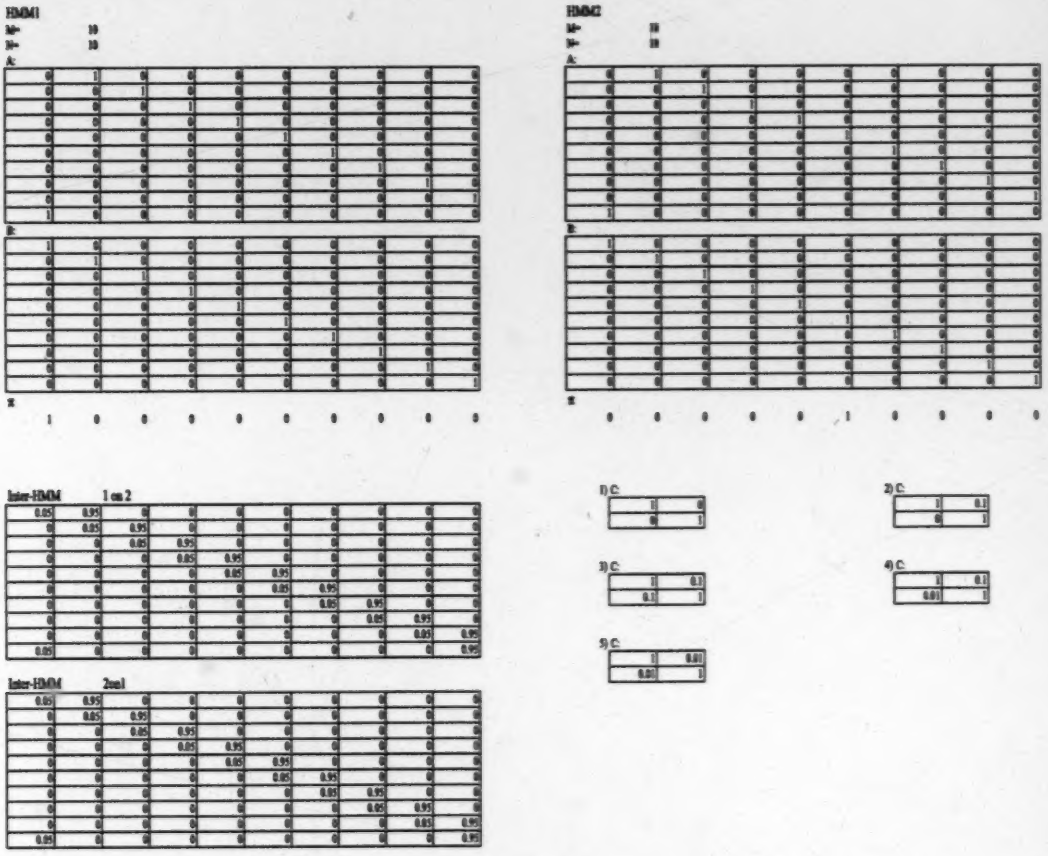

Table 6.3: Hypothetical model structure for a pair of deterministic coupled oscillators. The **A** structure generates a deterministic oscillation ($S_1 \rightarrow S_2 \ldots S_{10} \rightarrow S_1$).

was the maximum indeterminacy that retained an oscillatory pattern.

6.2.3 Model Performance

Model performance varied widely as a function of oscillator variability (i.e., noise), coupling strength and symmetry. In general however, across all oscillator configurations, very few combinations of parameters demonstrated the transition behaviour expected based on human experimentation. For example, in the deterministic oscillator exam-

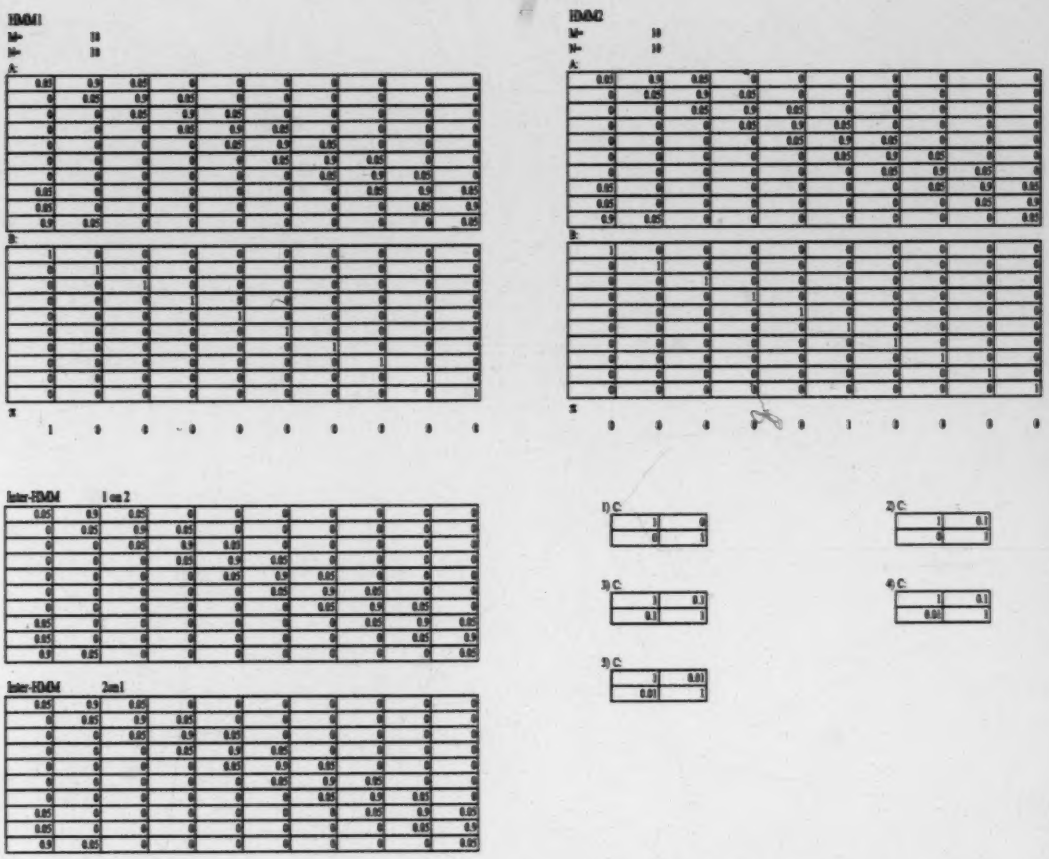

Table 6.4: Hypothetical model structure for a pair of pseudo-deterministic coupled oscillators. The **A** structure generates an oscillation ($S_1 \rightarrow S_2 \ldots S_{10} \rightarrow S_1$) with 5% noise.

ple, in all conditions but 2, transitions occurred immediately (within one oscillation). Although the transition was stable, and no return transition occurred, this result appears incongruent with human data. Conversely, the two conditions showed promising results; demonstrating appropriate behaviour were those models with a 1% coupling strength, either symmetric or unidirectional (see Figure 6.8, unidirectional shown). In the demonstration shown, a stable anti-phase relation was attained and maintained; based on random sampling of the gDBN, a transition occurred, resulting in a stable in-phase relation. While no variation exists, thus not replicating a decrease in phase

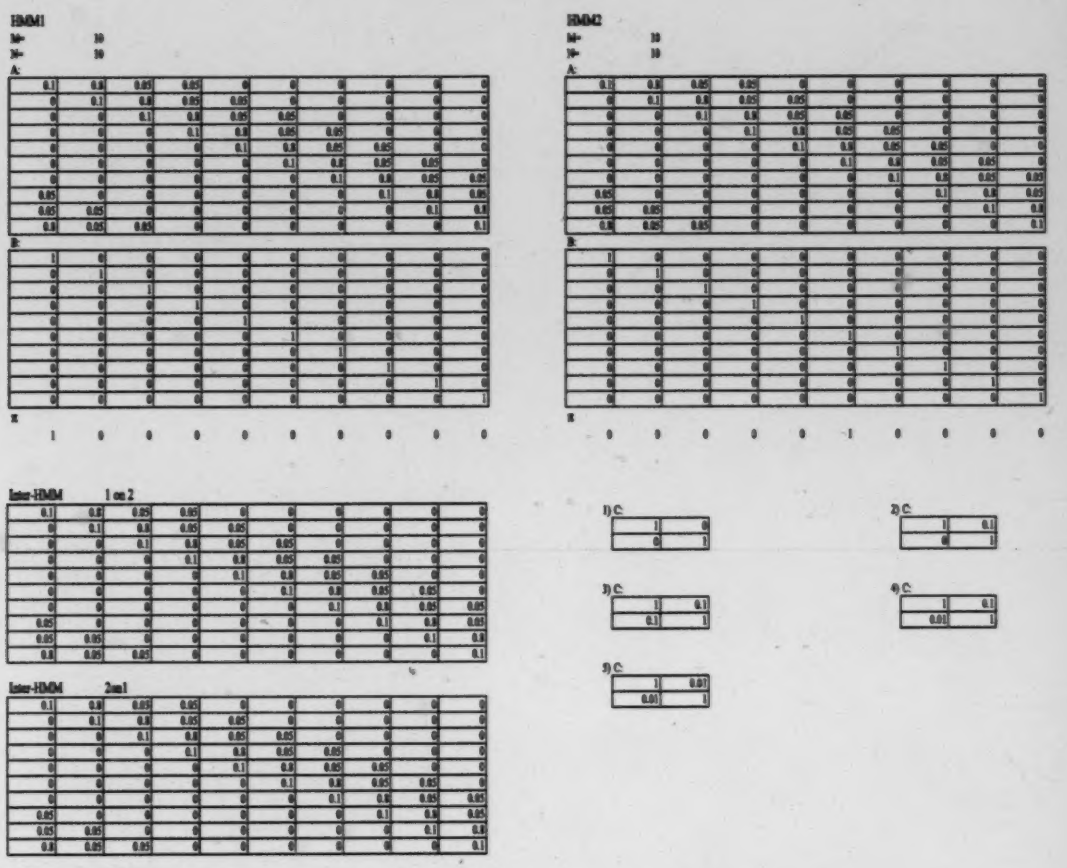

Table 6.5: Hypothetical model structure for a pair of noisy coupled oscillators. The **A** structure generates a stochastic oscillation from $(S_1 \rightarrow S_2 \ldots S_{10} \rightarrow S_1)$.

variability, the transition is certainly reflective of the behavioural predictions. Alternately, when the model was instantiated in an in-phase relation, no transition was present. In the second, semi-deterministic, simulation behaviour was much more reminiscent of human coordinated action; the initial phase relation was stable while still demonstrating some intra-phase variability. As the connection weight was increased, a much more consistent pattern of behaviours emerged. Under light coupling ($< 0.01\%$) the coordination pattern remained in-phase; as coupling strength was increased however, a slight shift away from an in-phase pattern was evident. Finally, as coupling

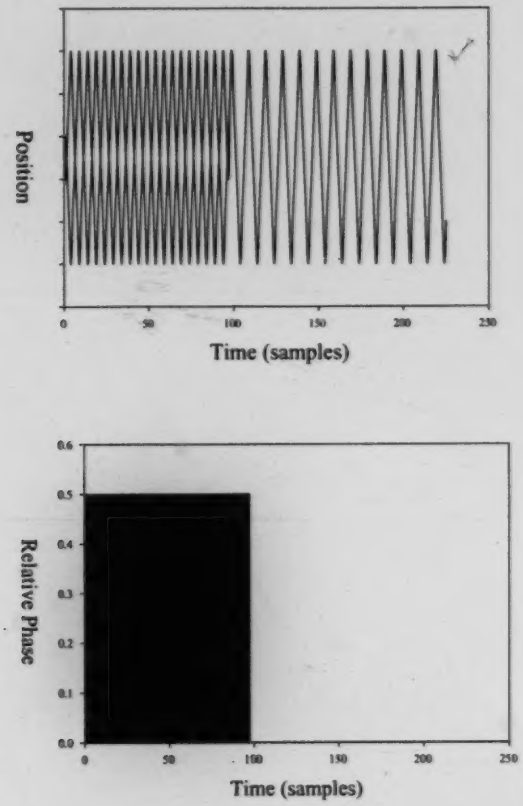

Figure 6.8: Cyclic behaviour beginning in an anti-phase configuration and undergoing a phase transition (deterministic oscillator conditions). Unidirectional coupling occurred with a relative strength of 1%.

strength reached 1%, phase transitions began to emerge. However, unlike the deterministic example, uni-directional connections demonstrated unstable phase transitions (i.e. transited back to anti-phase). In only one tested condition, asymmetric 1%/0.1% coupling (see Figure 6.9) did a reasonable transition occur; as coupling strength was increased beyond these values, random behaviour occurred with no discernible phasic pattern (see Figure 6.10). Similar to the deterministic oscillator model, when the semi-deterministic gDBN was started in an in-phase relation no transitions were present. However, unlike the previous example some variability remained present during in-

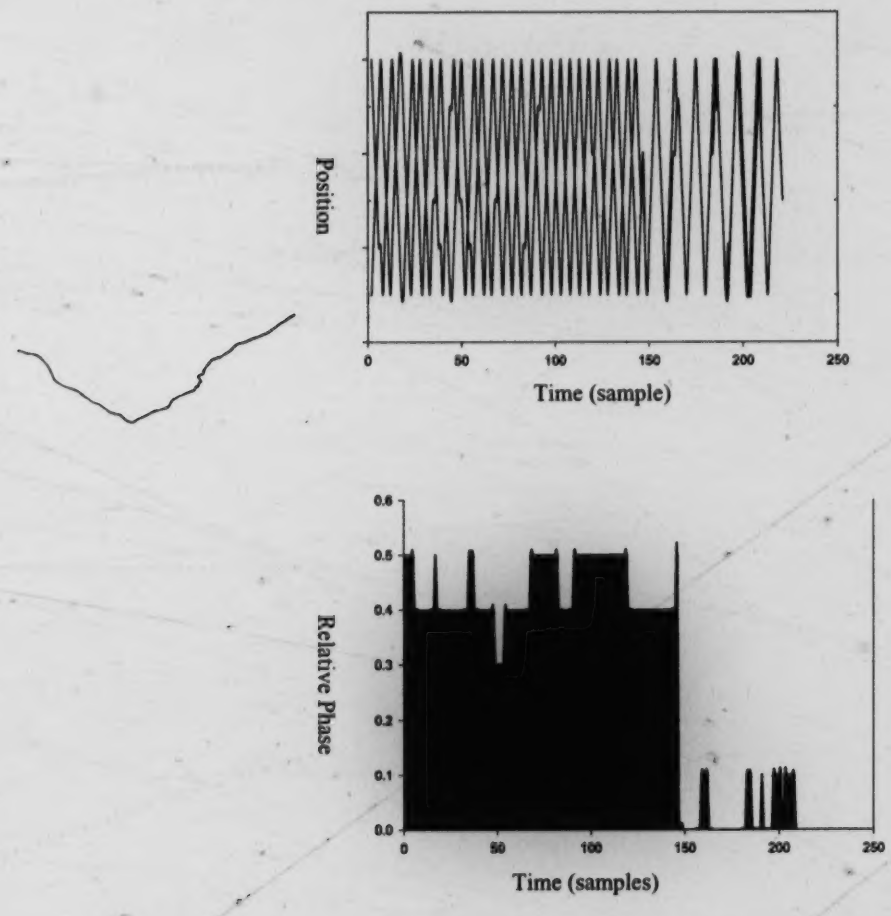

Figure 6.9: Variably cyclic behaviour beginning in an anti-phase configuration and undergoing a phase transition (semi-deterministic oscillator conditions). Asymmetric coupling of 1% and 0.1%.

phase coordination.

The final model condition (noisy oscillator) demonstrated random and aphasic in all coupling was environments. This result is contrary to expectations as the individual oscillators were able to maintain at least a pseudo-oscillatory pattern during independent operation, This observation places a theoretical bound on the noise level that a

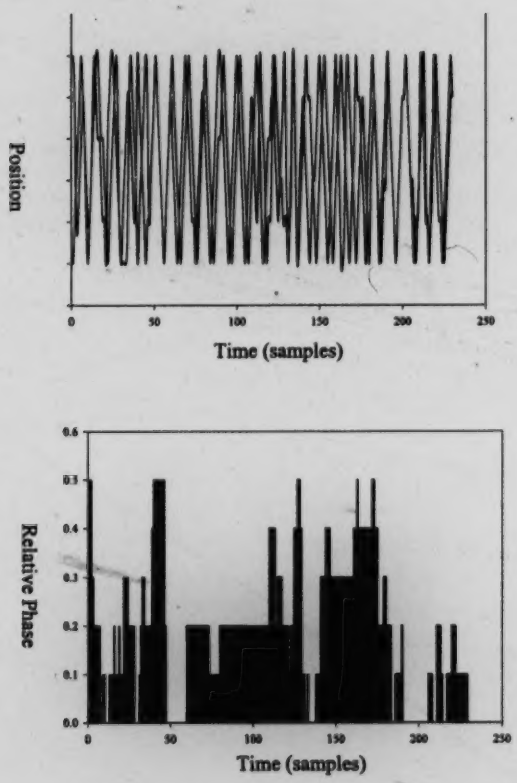

Figure 6.10: Acyclic behaviour beginning in an anti-phase start configuration (semi-deterministic oscillator conditions). Symmetric coupling conditions of 10%.

gDBN of this structure is able to represent (at least for cyclic behaviour). Perhaps this final demonstration provides a point of reference for the functional levels of 'noise' in a biological coordination system.

Up to this point, the discussion of the gDBN framework (with regard to continuous coordinated action) has focused on replicating the most overt predictions of the Haken-Kelso-Bunz (Haken et al. 1985) model. The results are very positive, however, the similarity of findings should not be unexpected as many of the internal features of the gDBN have parametric correlates within the HKB approach (see Equation 6.2, see also

Chapter 2 for a review).

$$V(\phi) = \delta - a\cos\phi - b\cos(2\phi) + \sqrt{Q}\zeta_t \tag{6.2}$$

Specifically, in parallel to Haken and colleagues' suggestion of a time dependent noise term, ζ_t, the current method integrates 'noise' into the dynamics of the system. The indeterminism is distributed across the super-A (causal A) matrix as well as in the expression of the states, **B**. Similarly, scaling is applied only to the noise term, $\sqrt{Q}$, and Ω is generalized to all state processes within the DBN. Finally, the 'detuning term', δ, is merely combined with the a and b terms for prescribing connection direction and strength. Thus, while the explicit description of the models vary, the compositions appear more congruent than not, suggesting the need for further formal investigations into their relations and the theoretical implications thereof – perhaps, in combination earlier results, even indicating a more parsimonious implementation of coupled oscillators toward discrete actions.

6.2.4 Conclusions

Overall, performance of the coupled oscillator simulations was positive, even given the simple coupled hidden Markov model case examined. Parameter configurations were demonstrated which reproduced phase transitions; moreover, simulations demonstrated pre-transition stability followed by transient instability with a transition followed by regaining stability at the in-phase relation. This replication in combination with a demonstration of hysteresis are consistent with models such as the HKB formulation; the gDBN may indeed be a generalization of coupled oscillator approaches. However, although some features were reproduced, many of the additional features accounted for by the HKB-model and contemporaries were not addressed. Further investigation needs to be made to examine features such as critical slowing and critical fluctuations.

Also with reference to coupled oscillator approaches, the gDBN formulation seems ideal

for examinations of features such as detuning, stochastic components, and coupling behaviour. As has been demonstrated by a Sternad and colleagues (Sternad, Amazeen & Turvey 1996) these features can be difficult to examine due to complex interactions and limitations in manipulation. Specifically, these authors suggest that detuning and coupling may covary. Comments are also made regarding the benefit of stochastic terms but no reference is made to their role. The gDBN method, in particular the ability to adjust features such as coupling strength and behaviour, allows parameters to be independently adjusted. Moreover, features such as oscillator eigenfrequency, although not examined here, can be manipulated explicitly to explore effects such as shifting of fixed point attractors away from 0 and π. Perhaps this approach would enable researchers to better classify and describe coupling – beyond synaptic, synergetic, or diffusive.

However, while this demonstration was beneficial for referencing coupled oscillator representations in a gDBN it can be criticized for merely "fitting the phenomena" instead of acquiring structure directly from observation; model parameters could just as easily have been set to display anti-phase stability. In the final experiment in this presentation of the gDBN method for prototyping action model, structure is "learned" directly from continuous participant data. Thus, unlike previous models of coordination the internal components of the model may be examined to better understand what happened instead of attempting to fit a function to the behaviour.

Chapter 7

On Learning the Internal Structure of a Biological Oscillator

7.1 Experiment 4: Modeling a Human Oscillator system

Over the last decade and a half, substantial research established the properties of the spatial and temporal constraints that exist for the execution of simultaneous, multi-limb movements. These constraints, while easy to demonstrate in a task such as "rubbing your tummy and patting your head", have generated what has been termed a paradigm shift in the study of the control of human movement: a transition toward examining human movement as a complex dynamical system whose behaviour can be better understood by studying relative phase instead of reaction time. As previous discussion has highlighted, current dynamical systems modeling practices successfully account for a wide variety of such continuous coordination tasks, both purely motor and perceptual-motor. However, although the abstract nature of the parameters have enabled some

flexibility in the accounts, it has limited the approach to descriptions of organization from that same level, making any formal model update from observed data very difficult. This final investigation examines the utility of generalized dynamical Bayesian networks, and the associated update procedures, for this type of parameter estimation problem.*

In addition to appraising the Baum Welch routine for capturing the temporal and stochastic features of a coordination task (cyclic pronation-supination), the pattern recognition properties of gDBNs are also investigated. Specifically, much of the basis for hypotheses regarding the complex features of coordination lie in the existence of non-linear behaviours (e.g., phase transitions) or the identification of attractor states in the dynamic 'coordination landscape' of the participant (e.g., phase scanning; Tuller & Kelso 1989). Common to both of these approaches is also a requirement for summarizing the action evolution within a collective variable (see Chapter 1). Consider a task, perhaps eye-hand coordination, where the coordinated effectors are incongruent with a relative phase calculation; a substantial portion of the eye behaviour is ballistic bursts (saccades) followed by fixation. Both previously mentioned assessment tools would be unable to classify the coordination in such a task. Moreover, even in tasks where relative phase (a collective variable) is present, some limb configurations may be impossible due to physical limitations (e.g., 90^o out of phase walking).

Given the success of Experiment 3 in reproducing some of the key behaviours in oscillatory coordination, a similar model formulation was used in this investigation. A two-node oscillator model was trained on continuous coordination participant data. A unique model was constructed for each of pre-transition, transition, and post-transition plateaus. By constructing models in this manner, the efficacy of various coupling properties could be examined for reproducing the recorded data. Moreover, data was collected for each 'oscillator' independently. Thus the coupling interactions could be

*Although some researcher may choose to map these estimated parameter values onto specific theses of control or biological organization, within a gDBN the values retain their statistical descriptive power regardless of translation.

estimated versus trying to extrapolate connection properties from collective variable behaviour. In addition, the updated models and associated internal transition likelihoods for each coordination stage can be tested for suitability across other behaviours (e.g., a model is constructed to represent stable anti-phase behaviour is applied to post-transition (in-phase) data). This cross comparison will allow the assessment of behaviour similarities – even when the overt behaviour may be quite disparate.

7.1.1 Methods

7.1.1.1 Participants

Three female student volunteered to participate in this investigation. All participants were naïve to the experimental hypotheses. Data collection and recruitment was completed according to the University of British Columbia guidelines for ethical treatment of human research subjects.

7.1.1.2 Task

Participants made continuous pronation and suppination actions with both their left and right wrists. Movements were made in each of two fashions. In-phase (IP) actions consisted of simultaneous pronation and suppination of both writsts. Anti-phase movements consisted of pronation with one wrist concurrent to suppination of the other. Complete cycles of movement were completed in synchrony with an auditory metronome.

While participants were instructed to maintain timing with the metronome, participants were also instructed to "do not resist" disruptions in the movement pattern. This instruction set entailed that subjects continue to move 'naturally', even if this required changing the limb coordination pattern (i.e transition).

7.1.1.3 Apparatus

Pronation-supination movements of the forearms were recorded using a pair of custom built levers (diameter: 2 cm, length: 18cm). The levers were mounted at shoulder width on a table-top directly in front of the seated participants. Angular displacement was monitored with the use of a rotational potentiometer (single turn, 0.25% linearity) attached coaxially with the center of rotation of the lever (see Figure 6.7 in previous Chapter). The signal from the potentiometer was sampled at 250 Hz with a 12 bit, 8 channel, A/D converter (National Instruments)

7.1.1.4 Procedure

Participants begin each session by familiarizing themselves with the inverted pendulum apparatus until confident with manipulation. Five (5) experimental trials were completed in both in-phase (IP) and anti-phase (AP) initial conditions. Participants alternated between the in-phase and anti-phase coordination pattern trials; starting phase condition was randomly assigned across subjects.

At the start of each trial, the metronome was set to a frequency of 1.0 Hz. The participants were to then attain the demonstrated frequency in the prescribed coordination (IP or AP). Throughout the duration of a trial, metronome frequency was scaled, in seven 0.25 Hz increment plateaus, from 1.0 Hz to 2.5 Hz. Each plateau lasted a total of 8 seconds. Scaling of the metronome frequency began upon indication of participant readiness.

7.1.1.5 Continuous Phase Angle & Relative Phase Estimates

Relative phase values for the pronation/supination task was calculated in a continuous fashion (Beek & Beek 1988). More explicitly, continuous phase angle (θ) of each limb

(potentiometer) was calculated for each sample (j) of angular displacement such that:

$$\theta_j = \arctan(\frac{\nu_j}{\Delta x_j}) \tag{7.1}$$

where ν is the velocity of the displacement ($\dot{x}$) divided by the mean frequency for the current plateau ($\bar{f}_n$).

$$\nu_j = \frac{\dot{x}}{\bar{f}_n} \tag{7.2}$$

Thus, *relative* phase is defined as difference between the continuous phase values for each hand.

$$\Phi_j = \theta_{j(left)} - \theta_{j(right)} \tag{7.3}$$

7.1.1.6 Data Summary and Clustering

Descriptive measures for each frequency plateau were calculated independently for each trial and subject. Only the final seven second of each plateau was used for calculations to ensure attainment of the criteria frequency.

Circular measures of central tendency of relative phase values were calculated according to Mardia (1972) (see also Burgess-Limerick, Abernathy & Neal 1991). These statistics were chosen over conventional values ($\bar{x}, \sigma$) due to computational errors encountered while describing circular distributions.[†] In summary, the average direction (vector) is used in place of arithmetic mean while uniformity is used in place of standard deviation. The circular variance (uniformity) is defined over a range of 0 to 1, however a transformation readily converts this to the continuous range of 0 to ∞ for inferential calculations.[‡].

[†]A good point to consider here is the problem of calculating an average direction from a selection of compass readings (from 0^o *to* 360^o): 7^o, 342^o, 350^o, and 16^o. The arithmetic mean of these values is 179 degrees with a standard deviation of 193. However, these numbers are unreasonable, more appropriate values would be in the order of $\bar{x} = 359, \sigma = 15$.

[‡]$s_t = (-2\log_e(1 - s_c))^{0.5}$ where s_t is the transformed variance and s_c uniformity

7.1.1.7 Model Construction and Assessment

Prior to the construction and training of the gDBN models, phase values for each limb were clustered in a ramp-fashion. The oscillatory amplitude was divided into 18 sequential observations representing 10 positions (see Figure 7.1). A total of 40 gDBN

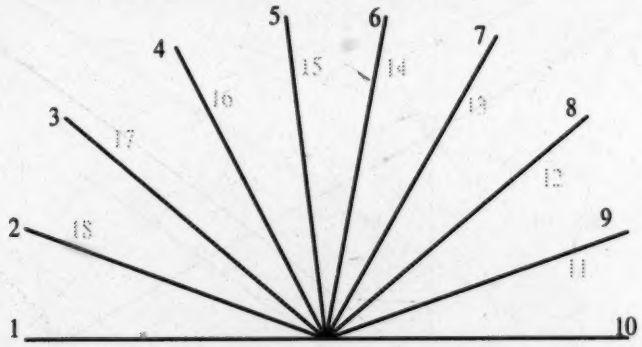

Figure 7.1: Schematic of discrete clustering method for a 10 position (18 observation) oscillator. Observation segments represent segments of an angular movement through 180°. Note that the #1 position and the #10 position are only encoded by one state as they represent movement reversals.

configurations were examined in this study. Specifically, a set of models was trained on each of anti-phase (AP), transition (TRAN), in-phase post transition (IPpost), and inphase (IP) data, grouped by movement frequency plateau. These divisions were established independently for each participant, based on continuous relative phase calculations and observed phase-shifts from the desired (anti-phase) relation to the in-phase relation. A classification of a transition frequency was made for plateaus where each participant made the majority of their transitions. AP and IPpost classifications were assigned to the neighbouring frequency plateau (i.e., for a TRAN plateau of 1.75 Hz, AP and IP post plateaus of 1.5 and 2.0 (respectively) would be selected) In the case of trials instantiated with an in-phase relation only one class of models was constructed. Within each class (AP, TRAN, IPpost, IP) 10 weighting configurations were examined. Specifically, models were trained on randomly selected class data (across subjects) with Ω connection values of (0%, 0%)(100%,100%), (10%,10%), (1%,1%), (100%,0%),

(0%,100%), (10%,0%), (0%,10%), (1%,0%), and (0%,1%).[§] Following training model performance was assessed relative to a novel data set (again randomly sampled from participant sequences) from all classes using a gViterbi evaluation of the most likely state sequence followed by a constrained MonteCarlo and MSD measure.[¶]

7.1.2 Discussion and Results

As is evident in Table 7.1, participants successfully attained their initial coordination goals. Moreover, by examination of the relative phase relations, in-phase destabilization is evident for each participant beyond a critical oscillation frequency. The graphs of such phase transitions also demonstrate the tendency for movement variability to decrease post transition (Figure 7.2). These data are consistent with the HKB model; when the coordination system is driven beyond a critical state, it self-organizes toward another "attractor" state. The demonstration of this non-linear shift in overt behaviour has been replicated many times and is therefore not the focus of this investigation. Of more interest was whether this data set could be used to train a series of gDBN models to replicate the participant behaviour based solely upon variations in coupling strength and symmetry – the variables hypothesized to scale as the perceptual-motor system is taxed. Beyond this hypothesis, there was also a goal of statistically summarizing the system behaviour at stable and unstable points; regardless of the success of the coupling hypothesis this result would demonstrate the general feature of HMMs and gDBNs for prototyping a dynamical system based on sequential observations.

Pertaining to the prototyping and replication goal of this investigation, several features

[§]Percentage connection values are represented in the Ω matrix within the gDBN. Values displayed here may be read as: (1) left hand influence on right; and right hand on left. Percentage scale is the magnitude of connection relative to the within node (hand) connection value – in this examination this value was always 1 (100%).

[¶]The included data tables must be examined in an orderly fashion. The log(p—L) values represent the likelihood that a sequence (the gViterbi generated one) would happen. Thus, regardless of the the other number performances, if this number is small (i.e., highly negative) then the model is inneffective at data representation. If the log probability is reasonable however, then the Hamming score and MSD can be examined to determine the degree of fit of this sequence based on predicted observations

Plateau Freq.		Starting Coordination Pattern: In-Phase Subj 1	In-Phase Subj 2	In-Phase Subj 3	Anti-Phase Subj 1	Anti-Phase Subj 2	Anti-Phase Subj 3
	1	-5.68751	-7.38311	6.25457	173.919	173.3155	-175.412
	2	-0.70834	-5.28101	2.707342	169.968	170.8407	-174.223
0.75 Hz	3	-5.98288	-3.74625	1.639168	172.1309	170.7757	-175.442
	4	-2.09509	-6.32333	4.675634	173.7061	172.0178	-172.902
	5	-6.87878	-4.48329	3.178954	171.901	167.7468	-175.031
	1	-4.36852	-15.3893	5.661889	172.9944	161.0633	-172.359
	2	-5.93042	-10.4959	3.624753	167.5246	162.0008	-174.567
1.0 Hz.	3	-4.44619	-9.12365	-0.12298	170.4462	168.9609	-174.647
	4	-3.16621	-9.6005	2.19609	171.613	166.1378	-172.347
	5	-5.1224	-8.09879	3.405997	170.6817	163.8535	-169.954
	1	-7.44776	-14.9543	4.815001	169.2301	155.367	-172.286
	2	-7.92359	-9.38902	-0.16047	165.5243	159.2785	-168.969
1.25 Hz.	3	-8.28625	-16.7008	0.569401	164.9327	160.2011	-174.943
	4	-5.32606	-10.5618	0.77156	166.2815	158.8451	-166.888
	5	-8.48037	-9.71245	-1.94162	161.9432	159.619	-172.224
	1	-8.72017	-14.9866	5.667345	167.6348	155.5344	178.7646
	2	-9.77449	-13.9496	-4.6449	162.7758	149.4181	-176.058
1.5 Hz.	3	-9.04643	-15.0031	-0.48608	162.2273	154.8856	175.1265
	4	-6.47117	-7.46164	1.447535	168.3811	147.7977	164.1789
	5	-9.67531	-15.9149	1.033633	161.1832	146.4639	171.9303
	1	-8.81558	-15.5987	1.684873	178.1005	139.6547	176.4482
	2	-13.923	-26.0418	2.237374	166.9159	130.9441	65.0649
1.75 Hz.	3	-12.4435	-14.7905	-2.64522	167.3173	155.5336	-6.01401
	4	-6.78448	-15.6614	2.093533	-167.299	147.4054	-2.98145
	5	-11.7124	-17.0919	-2.61749	-172.927	138.6921	18.93486
	1	-5.54158	-13.0426	-5.46019	-18.4796	126.0551	16.89333
	2	-11.6138	-26.697	-6.80958	-1.29799	125.1022	0.613117
2.0 Hz.	3	-10.2253	-20.1432	-4.13363	-179.273	127.5976	-7.86285
	4	-6.42959	-15.701	4.132981	-6.90305	114.9509	-11.9248
	5	-12.2721	-14.7638	3.926155	0.548553	134.143	-6.57041
	1	-6.47375	-11.3516	2.082441	-0.68945	165.5397	3.424451
	2	-7.6754	-24.8125	-3.40209	-3.34419	106.9974	0.064112
2.25Hz.	3	-10.7416	-12.7207	-10.3684	0.25056	117.0535	-5.19089
	4	-8.86305	-9.53002	-5.79766	-3.42637	115.7761	-13.3791
	5	-8.21011	-8.53894	0.319829	-6.47561	90.19324	-5.85981
	1	-2.25454	-1.29391	-4.90073	3.059222	84.45869	-1.94717
	2	-8.85664	-12.0406	-2.92574	-0.34975	119.0302	-10.0016
2.5	3	-8.66613	-7.85433	-13.2347	-3.72669	118.7147	-13.4551
	4	-8.11805	-2.52056	-6.63212	-1.07607	58.99432	-7.97849
	5	-4.91588	1.728021	-9.34883	0.754367	69.99004	-14.6234

Table 7.1: Mean relative phase values for participants across all start phase relations and frequency plateaus. Summary statistics calculated according Mardia 1972.

		AP			TRAN			IP post			IP		
L->R	R->L	H Distance	MSD	log10(P(O|L))	H Distance	MSD	log(P(O|L))	H Distance	MSD	log(P(O|L))	H Distance	MSD	log(P(O|L))
%	%												
0.00	0.00	1.00	0.00	-188.14	0.07	8.47	-	0.10	7.83	-	0.26	5.94	-
1.00	0.00	0.91	0.10	-295.36	0.47	4.47	-	0.45	4.07	-	0.49	3.84	-
0.00	1.00	0.95	0.10	-201.87	0.82	0.82	-	0.75	0.55	-	0.79	0.40	-455.00
1.00	1.00	0.98	0.02	-223.41	0.99	0.01	-302.78	0.98	0.02	-459.26	0.98	0.02	-388.47
10.00	0.00	0.94	0.07	-269.34	0.48	4.45	-	0.47	4.04	-	0.56	3.13	-
0.00	10.00	0.95	0.10	-202.97	0.82	0.82	-	0.77	0.48	-504.68	0.81	0.39	-427.40
10.00	10.00	0.96	0.04	-253.60	0.89	0.12	-303.41	0.94	0.06	-396.08	0.94	0.06	-375.58
100.00	0.00	0.94	0.06	-185.45	0.48	0.45	-	0.52	3.27	-	0.57	3.01	-
0.00	100.00	0.91	0.17	-216.89	0.80	0.86	-	0.76	0.46	-457.95	0.78	0.42	-417.12
100.00	100.00	0.88	0.12	-299.71	0.88	0.19	-417.54	0.83	0.39	-447.78	0.84	0.34	-428.31

Table 7.2: AP models versus each of TRAN,IP,IPpost, and AP data. The left-most column depicts the coupling behaviour in % maximum. H Distance and MSD account for the degree of correspondence between traces using a Hamming distance measure and mean distance respectively. Independently of the distance measures, log10(p(O—L)) represents the likelihood of the predicted sequence; if this measure is either very small or "-" then then the simulated observation sequence is non-conforming (regardless of the correspondence measure).

of model performance are important to note (see Tables 7.2 to 7.5; see also Appendix B for a model example). In general, the gViterbi solution was exceptional at predicting the recorded observation sequence (e.g., Table 7.2; 0%,0%). Measures of observation correspondence often reached values beyond 0.9 (Hamming score) and < 1.0 discrete observations (MSD); this reflects data being represented accurately over 90% of the time and, when errors are made, they occur on average only one 'step' away from the recorded value. Also, model performance varied as a function of both the independent measures – coupling and plateau frequency (i.e., IP, IPpost etc) – showing sensitivity to the experimental manipulation. As should be expected, scores were almost invariably superior for similar trials. This demonstration of recognition, while important, should be evaluated with caution. In order for a model to be useful it must show recognition of 'itself' as well as discrimination between similar and dissimilar data sets. This second feature should also be considered within the framework of a phase-transition paradigm. Specifically, if a coupling-strength variation model is to be considered successful at representing the the state of a control variable it should:

1: Recognize similar systems

AP models should be likely when evaluating any anti-phase data set

IPpost model should recognize all self-similar conditions (IP & IPpost)

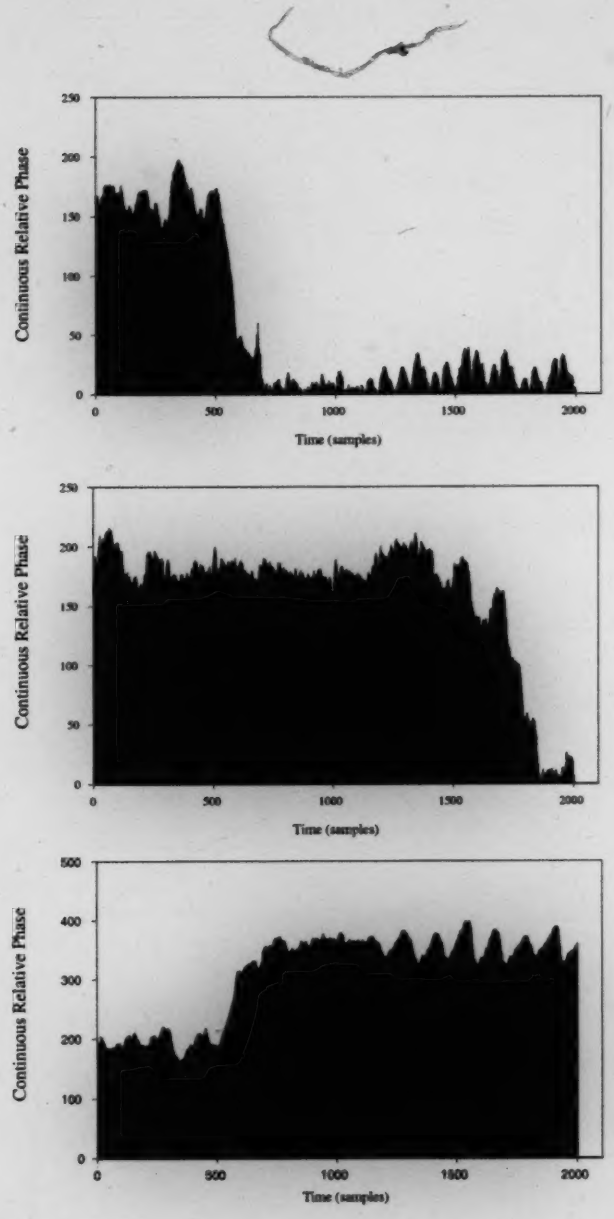

Figure 7.2: Example phase transitions from participant data. *Upper Panel*: Data from trial 3 of participant #1 (2.0 Hz. plateau). *Middle Panel*: Data from trial 1 of participant #2 in the (1.5 Hz. plateau)*Bottom Panel*: Data from trial 5 of participant #3 in the (1.5 Hz. plateau).

		AP			TRAN			IP post			IP		
L -> R	R -> L	H Distance	MSD	log10(P(O|L))	H Distance	MSD	log(P(O|L))	H Distance	MSD	log(P(O|L))	H Distance	MSD	log(P(O|L))
%	%												
0.00	0.00	0.11	6.95	-	0.50	3.17	-	0.07	7.13	-	0.11	6.75	-
1.00	0.00	0.33	5.48	-	0.75	1.41	-495.50	0.28	5.85	-	0.31	5.55	-
0.00	1.00	0.19	6.25	-	0.57	2.65	-420.51	0.13	5.97	-	0.16	5.72	-
1.00	1.00	0.62	1.85	-575.34	0.62	2.21	-	0.63	2.03	-605.82	0.62	2.08	-609.95
10.00	0.00	0.33	5.49	-	0.75	1.45	-494.46	0.28	5.87	-	0.31	5.57	-
0.00	10.00	0.19	6.26	-	0.57	2.64	-420.70	0.13	5.98	-	0.16	5.72	-
10.00	10.00	0.63	1.83	-572.99	0.63	2.21	-	0.64	2.00	-606.17	0.62	2.11	-
100.00	0.00	0.33	5.53	-	0.70	1.81	-189.48	0.28	5.87	-	0.31	5.58	-
0.00	100.00	0.43	2.83	-470.79	0.57	2.64	-420.55	0.31	.	-429.28	0.28	4.11	-448.41
100.00	100.00	0.61	2.88	-318.24	0.64	2.21	-235.50	0.60	3.77	-285.48	0.62	1.81	-264.09

Table 7.3: TRAN models versus each of TRAN,IP,IPpost, and AP data. The left-most column depicts the coupling behaviour in % maximum. H Distance and MSD account for the degree of correspondence between traces using a Hamming distance measure and mean distance respectively. Independently of the distance measures, log10(p(O—L)) represents the likelihood of the predicted sequence; if this measure is either very small or "-" then then the simulated observation sequence is non-conforming (regardless of the correspondence measure).

2: Discriminate incongruent data

AP models should discriminate between AP and IP

IP model should discriminate between IP and AP

3: Be *somewhat* consistent with neighbouring transition data

Given these restrictions, some tendencies appear. In particular, only the symmetric coupling parameterized models accurately replicate behaviour under anti-phase conditions. Conversely, only asymmetric models prototype in-phase behaviours – either IP or IPpost (see Tables 7.2 and 7.5). Contrary to predictions however, no models successfully encapsulate all requirements. Attempts either were unable to descriminate between inconsistent modes (e.g., AP & IP; Table 7.2, 100%, 100%) or were unable to recognize self-similarity (e.g., Table 7.4, 0%, 0%).

While there are many possible explanations for these null results, two proffer testable hypotheses. The first account is that similar behaviour may be generated by divergent control parameters. For example, IP and IPpost may be at different, and likely not 'critical', positions on the control state spectrum. While this could arise from several conditions (e.g., instructions or initial conditions), if this were the case model agreement between generated parameters would not be expected. The second, some-

		AP			TRAN			IP post			IP		
L->R	R->L	H Distance	MSD	log10(P(O|L))	H Distance	MSD	log(P(O|L))	H Distance	MSD	log(P(O|L))	H Distance	MSD	log(P(O|L))
%	%												
0.00	0.00	0.15	6.99	-	0.38	4.16	-	0.87	0.24	-200.66	0.46	4.16	-
1.00	0.00	0.51	3.36	-	0.80	0.25	-609.88	0.87	0.24	-199.33	0.85	0.33	-241.02
0.00	1.00	0.47	3.82	-	0.46	4.05	-	0.87	0.22	-210.89	0.84	0.32	-250.38
1.00	1.00	0.79	0.25	-512.45	0.86	0.17	-488.76	0.86	0.23	-211.92	0.85	0.22	-232.72
10.00	0.00	0.63	0.62	-701.15	0.79	0.26	-525.16	0.85	0.25	-201.62	0.83	0.34	-234.57
0.00	10.00	0.47	3.51	-	0.47	4.05	-	0.87	0.22	-215.53	0.84	0.32	-237.13
10.00	10.00	0.80	0.23	-455.07	0.87	0.16	-434.22	0.87	0.27	-223.26	0.85	0.19	-240.37
100.00	0.00	0.51	3.46	-	0.74	0.80	-338.55	0.84	0.29	-205.02	0.82	0.37	-223.19
0.00	100.00	0.47	3.51	-	0.46	4.05	-	0.86	0.23	-191.86	0.84	0.32	-230.85
100.00	100.00	0.89	0.11	-399.57	0.91	0.10	-338.28	0.89	0.11	-264.16	0.90	0.11	-287.35

Table 7.4: IPpost models versus each of TRAN,IP,IPpost, and AP data. The left-most column depicts the coupling behaviour in % maximum. H Distance and MSD account for the degree of correspondence between traces using a Hamming distance measure and mean distance respectively. Independently of the distance measures, log10(p(O—L)) represents the likelihood of the predicted sequence; if this measure is either very small or "-" then then the simulated observation sequence is non-conforming (regardless of the correspondence measure).

		AP			TRAN			IP post			IP		
L->R	R->L	H Distance	MSD	log10(P(O|L))	H Distance	MSD	log(P(O|L))	H Distance	MSD	log(P(O|L))	H Distance	MSD	log(P(O|L))
%	%												
0.00	0.00	0.11	6.81	-	0.05	7.43	-	0.22	5.35	-	0.47	3.08	-
1.00	0.00	0.55	3.87	-	0.48	4.45	-	0.67	2.41	-	0.91	0.25	-230.96
0.00	1.00	0.29	4.98	-	0.28	5.22	-	0.75	0.79	-250.64	0.77	0.75	-219.63
1.00	1.00	0.86	0.22	-566.43	0.86	0.27	-395.43	0.89	0.26	-243.68	0.89	0.27	-230.91
10.00	0.00	0.55	3.87	-	0.51	4.01	-	0.67	2.42	-	0.91	0.25	-232.56
0.00	10.00	0.29	5.16	-	0.56	2.04	-	0.75	0.79	-246.12	0.77	0.72	-220.16
10.00	10.00	0.85	0.20	-435.23	0.86	0.22	-391.28	0.90	0.17	-247.18	0.90	0.22	-235.66
100.00	0.00	0.55	3.87	-	0.51	4.10	-	0.81	1.07	-	0.92	0.25	185.06
0.00	100.00	0.56	1.77	-579.31	0.59	2.00	-	0.75	0.80	-247.31	0.75	0.76	-220.58
100.00	100.00	0.76	0.77	-399.87	0.72	0.73	-312.66	0.77	0.56	-289.60	0.77	0.54	-258.01

Table 7.5: IP models versus each of TRAN,IP,IPpost, and AP data. The left-most column depicts the coupling behaviour in % maximum. H Distance and MSD account for the degree of correspondence between traces using a Hamming distance measure and mean distance respectively. Independently of the distance measures, log10(p(O—L)) represents the likelihood of the predicted sequence; if this measure is either very small or "-" then then the simulated observation sequence is non-conforming (regardless of the correspondence measure).

what more tenable explanation suggests that the model parameters were insufficient to capture all the subtleties of the system. Regardless of the basis, further examination of the underlying parameters of simple coordination is needed before any reasonable conclusion can be reached.

7.1.3 Conclusions

Overall, the results of this final investigation again emphasized the utility of the generalized dynamical Bayesian network for applications in modeling human movement.

Even though some problems have been raised regarding the performance of models within the scope of the current coupling hypotheses, the approach cannot be summarized as anything else but successful. Continuous, complex, non-linear behaviours were accurately represented by the acquired models, sometimes accounting for $> 95\%$ of the data variance! Moreover, recognition of pre- and post-transition behaviours was accomplished even given the somewhat ambiguous classifications; participants may have made transitions earlier, later, or not at all – classification was made with reference to their modal performance.

This method has demonstrated great representational efficacy, along with appropriate and robust estimation methods, suggesting its application to more focused investigation of internal motor parameters in movement. For example, one might examine the role of control parameters on not only the direction and strength of coupling, but also the 'shape' of coupling. Current models leave little flexibility for changes in the fundamental properties other than coupling (b/a); the expression of coupling, whether due to the fundamental eigenfrequencies of the oscillators or neuromotor noise, is not testable under the HKB framework. Using gDBN methods, not only the static values of these paremeters are examinable (Appendix B), but also their evolution. Models trained on consecutive frequency requirements can be compared at a sub-matrix level, examining the dynamics of features such as oscillator noise and perhaps symmetry as a function of the control parameter.

Regardless of the achievements of the gDBN models for reproducing behaviour, the final point of this discussion should be the major shortfall of this approach, that is, the reliance on a temporally *complete* search of parameters during evaluation (i.e., gViterbi). Although this approach does adequately test the representational power of any model, it performs this test in an manner inconsistent with biology. For example, it is unreasonable to propose that a control system assess its entire performance (e.g., 30 sec) via feedforward processes prior to movement. This evaluation would be the only way that a Viterbi search could control action. Perhaps, this is an inappropriate

expectation; perhaps instead a procedure such as this is used at movement completion – to assess for the purposes of learning.

Chapter 8

On Reaching for a Cup of Coffee: Revisited

Having had an extensive exposure to the current approaches to the study of human movement, as well as the framework posed here, let us reconsider the simple task of picking up a cup of coffee and how one can study it. This movement problem requires the rapid inclusion, via computation or otherwise, of a large number of perceptual and motor factors; a problem which has traditionally required an "all-knowing" homunculus to act as the puppet master. The focus of this dissertation has been the development, discussion and testing of Bayesian process models (i.e., HMMs and gDBNs) for prototyping human actions. The framework has demonstrated notable success at replicating a variety of tasks, distinguishing it from many of the current modeling approaches in the literature on several fronts. Likely the most important feature of the gDBN/HMM method is its uses and inclusion of variability into its structure; unlike deterministic models of action, the indeterminacy of the motor control system is accounted for implicitly within the internal parameters of the model. A property that should also be considered is the inherent dynamic behaviour of this framework, while retaining a process-oriented or modular format. In concert with transparent, flexible and proven

optimization, this dynamic modularity separates the gDBN/HMM methods of control representation from competitors such as ANNs and dynamical pattern theory, making it a viable choice for researches wishing to prototype action. However, although significant success has been demonstrated, the immaturity of the approach demands further refinements for movement science. Above all, the model is very *general*, being applicable to any complex motor task while still providing a principled, optimal parameter estimation.

8.1 Extensions

8.1.1 Improving Ω Through gBaum Welch

As was briefly discussed in Chapter 5, there are many ways to envision the update procedures associated with generalized Dynamical Bayesian networks. The format discussed and implemented in this dissertation utilized a "mass action" approach to combining evidence from multiple nodes in determining the behaviour of each HMM independently. While this contributor approach corresponds subjectively to other approaches to machine learning (e.g., artificial neural networks) and Bayesian networks (e.g, Ghahramani & Jordan 1997) it is unable to replicate some of the features of a physical movement system. For example, there should be a mechanism within such a model do dissallow individual states; in the current implementation, the context provided by the weighting matrix and inter-HMMs can only increase the likelihood of an event – not prevent it.

Two solutions present themselves to extend the mass action model for this purpose. In the method described by equation 8.1, only a slight change has been made; the range of values permitted for the weighting term has been extended below zero. The implications of this change however allows the effect of a node on the behaviour of another to be inhibitory; an effect consistent with other mass action approaches (c.f., ANNs) and

some biological behaviours (e.g., presynaptic inhibition, inhibition of return). Although promising for extending the current method, this amendment is still unable to evoke inhibition of individual state transitions.

$$q^i_{t+1} \propto \left[\left(\sum_{i=j,u} \omega^{ii}(u) a^{ii}_{uv} \right) + \left(\sum_{j \neq i,y} \omega^{ji}(y) a^{ji}_{yv} \right) \right] ; \; For -1 \leq \omega^{ji} \leq 1 \tag{8.1}$$

The second proposed change is somewhat more fundamental; quite simply, the combination of inter-nodal terms is done in a multiplicative fashion (Equation 8.2) instead of an additive one. The effect of this change however can be demonstrated by simply considering a scenario where the inter-HMM (off-diagonal of the causal **A**) contains a zero. The presence of a zero will make the occurrence of that transition *impossible*.

$$q^i_{t+1} \propto \left[\left(\sum_{i=j,u} \omega^{ii}(u) a^{ii}_{uv} \right) \times \left(\prod_{j \neq i,y} \omega^{ji}(y) a^{ji}_{yv} \right) \right] \tag{8.2}$$

While this second method appears to have many benefits, such as inhibitory control at a state level, there are problems. For example, consider the case where all transitions are impossible due to inter-connection restrictions. What is the expected behaviour of the system? The more substantial concern however is the implication for Expectation Maximization; although the execution behaviour of the model seems reasonable, the effect of many zeroed terms on update and recursion may be more difficult.

A final consideration for improvements to gDBN update and production involves generating state expressions; again there are two extensions. The first involves the observations being determined base on a continuous probability density as opposed discrete observations. While the basic expression for the gDBN algorithms would remain much unchanged with this approach (e.g. $o^j \propto p(o^j|x^j)$ the implications for model update might be extensive. In addition, the variable type permissible for legitimate observations would be increased to continuous, ratio data. The second adjustment requires a substantial algorithmic and logical extension – having the behaviour of one node influence the *expression* of another (i.e., $o^j \propto p(o|x^{ji})$; $\forall i, j$). While some-

what against the tradition of HMMs, this approach is definitely reminiscent of the the theoretical formulation proposed by Berkinblit for Tsetlin automata (Berkinblit et al. 1986) and, by extension, proponents of equilibrium point hypotheses for action (; Latash 1993, Feldman 1986) the components of the system would be linked based on their common output.

8.1.2 A 'Local' gViterbi Controller

As was alluded to in the final section of Chapter 7, the gViterbi procedure makes some unreasonable assumptions if it is to be implemented as a real-time controller. Specifically, the algorithm requires knowledge of the entire sequence to be able to determine the optimal state sequence for it to arise from. The obvious solution, and the one suggested, is to restrict the extent of the optimization a local search – a "working motor memory" model. As such, for $t - n$ *to* t, recurse as before, such that Φ^i is vector of the likelihood of each state given the observation. This local recursion would occur across the recent 'memory' of the system – effectively estimating the control process which just occurred or perhaps evaluating the best method of attaining a desired outcome. Again, while the implementation seems straight forward, the implications of this change raises questions regarding issues like optimal memory size and local versus global optimization.

Perhaps the best way to consider this adaptation of Viterbi is to frame it within a classic information processing model for movement production. For example, during the "movement planning" portion of generation a feedforward expectation of the action would be produced. While the exact temporal extent of this event is debatable, one might expect this duration to be at least long enough in duration to allow visual feedback (i.e., the window size of gViterbi). The goal output of the system could be assessed over this window and the optimal set of control states selected. During online controlling, the same process can occur in a feedback fashion, determining the events

which occurred recently (this allows for noise in the transmission and production of the commanded action). However, before becoming embroiled in a debate of the online control properties of gViterbi, it should be noted that many of the features it provides can be augmented or even replaced with a complete model design, improved connection implementation (see previous), and perhaps even increased system order (i.e., t-n to t).

8.2 Concluding Remarks

A number of problems for movement scientists have been discussed, a number of possible solutions presented. However, in light of this dissertation and the information it contains, where are we left? A debate still rages between proponents of information processing researchers and those who profess a dynamical systems approach; one side stresses the importance of understanding the *processes* of human control while the other consideres the emergent behaviour, suggesting that the component interact in a complex even chaotic fashion better examined at a collective level. Perhaps that is where this research belongs: bridging the "theoretical chasm" (Newell 2001) that has appeared in the motor control literature. This joining is important if ever a comprehensive understanding – via modeling or experimentation – of the action-perception is going to be gained.

Bibliography

Abrams, R., Meyer, D. & Kornblum, S. (1989). Speed and accuracy of saccadic eye movements: Characteristics of impulse variability in the oculomotor system, *Journal of Experimental Psychology: Human Perception and Performance* **15**: 529–543.

Asatryan, D. G. & Feldman, A. (1965). Functional tuning of the nervous system with control of movments or maintenance of a steady posture. I. mechanographic analysis of the work of the limb on execution of a postural task, *Biophysics* **10**: 925–935.

Beek, P. & Beek, W. (1988). Tools for constructing dynamical models of rhythmic movement, *Human Movement Science* **7**: 301–342.

Beek, P., Peper, C. & Stegeman, D. (1995). Dynamical models of movement coordination, *Human Movement Science* **14**: 573–608.

Berkinblit, M., Feldman, A. & Fukson, O. (1986). Adaptability of innate motor patterns and motor control mechanisms, *Biofizika* **31**(1): 128–138.

Bernstein, N. (1967). *The coordination and regulation of movement*, Pergamon, New York, NY.

Binsted, G., Cullen, J. & Elliott, D. (1998). Manual assymetries in goal-directed movement: Examination of the motor output hypothesis, *Brain and Cognition* **38**: 102–110.

Bizzi, E. (1980). Central and peripheral mechanisms in motor control, *in* G. Stelmach & J. Requin (eds), *Tutorials in motor behavior*, North-Holland, Amsterdam, chapter 7, pp. 131–143.

Bizzi, E., Accornero, N., Chapple, W. & Hogan, N. (1982). Arm trajectory formation in monkeys, *Experimental Brain Research* **46**: 139–143.

Bizzi, E., Ivalidi, F. M. & Giszter, S. (1991). Computations underlying the execution of movement: A biological perspective, *Science* **253**: 287–291.

Bizzi, E., Polit, A. & Morasso, P. (1976). Mechanisms underlying achievement of final head position, *Journal of Neurophysiology* **39**: 435–444.

Bootsma, R. J. & Mottet, D. (1997). Dynamics of trajectory formation and speed/accuracy trade-offs, *Behavioral and Brain Sciences* **20**: 303–304.

Bose, C. B. & Kuo, S. (1994). Connected and degraded text recognition using hidden Markov model, *Pattern Recognition* **27**(10): 1345–1363.

Boulic, R. (1999). Human motion capture driven by orientation measurements, *Presence: Teleoperations & virtual Environments* **8**(2): 187–203.

Brand, M. (1983). Coupled hidden Markov models for modeling interacting processes, *Learning and Common Sense Technical Report 405*, MIT Media Lab, Perceptual Computing.

Brand, M., Oliver, N. & Pentland, A. (1997). Coupled hidden Markov models for complex action recognition, *Proceedings of the IEEE: Computer Vision and Pattern Recognition*, San Juan, Puerto Rico.

Bui, H., Venkatesh, S. & West, G. (2001). Tracking and surveillance in wide-area spatial environments, *International Journal of Pattern Recognition and Artificial Intelligence* **15**(1): 92–107.

Bullock, D., Bongers, R. M., Lankhorst, M. & Beek, P. J. (1999). A vector-integration-to-endpoint model for performance of viapoint movements, *Neural Networks* **12**: 1–29.

Bullock, D., Cisek, P. & Grossberg, S. (1998). Cortical networks for control of voluntary movements under variable force conditions, *Cerebral Cortex* **8**: 48–62.

Bullock, D. & Grossberg, S. (1988). Neural dynamics of planned arm movements: Emergent invariants and speed-accuracy properties during trajectory formation, *Psychological Review* **95**: 49–90.

Bullock, D. & Grossberg, S. (88). The vite model: A neural command circuit for generating arm and articulator trajectories, *in* J. Kelso, A. Mandell & M. Schlesinger (eds), *Dynamic patterns in complex systems*, World Scientific Publishers, Singapore, pp. 305–326.

Burgess-Limerick, R., Abernathy, B. & Neal, R. (1991). Note: statistical problem in testing invariance of movement using the phase plane model, *Journal of Motor Behaviour* **23**: 30–303.

Byblow, W., Bysouth-Young, D., Summers, J. & Carson, R. (1998). Performance asymmetries and coupling dynamics in the acquisition of multifrequency bimanual coordination, *Psychological Research* **61**: 56–70.

Caelli, T., Binsted, G. & Frankland, J. (in submission). Encoding and recognizing actions using differential geometry and coupled hidden Markov models. *Pattern Recognition*, 38p.

Caelli, T., Guan, L. & Wen, W. (1999). Modularity in neural computing, *Proceedings of the IEEE* **87**(9): 1497–1518.

Caelli, T., McCabe, A., Binsted, G. & Williamson, R. (2000). Spatially invariant representation of human actions. Manuscript in preparation.

Carlton, L. (1981). Processing visual feedback information for movement control, *Journal of Experimental Psychology: Human Perception and Performance* **7**: 1019–1030.

Carlton, L. (1992). Visual processing time and the control of movement, *in* L. Proteau & D. Elliott (eds), *Vision and motor control*, North Holland, Amsterdam.

Carlton, L. G. & Liu, Y.-T. (1997). Speed/accuracy relations: The kinematic link and predictions for rapid timing tasks, *Behavioral and Brain Sciences* **20**: 304.

Carlton, L. & Newell, K. (1993). Force variability and characteristics of force production, *in* K. Newell & D. Corcos (eds), *Variability and motor control*, Human Kinetics Publishers, Champaign, IL, pp. 15–37.

Carnahan, H. & Marteniuk, R. G. (1991). The temporal oraganization of hand, eye, and head movements during reaching and grasping, *Journal of Motor Behavior* **23**: 109–119.

Carnahan, H. & Marteniuk, R. G. (1994). Hand, eye, and head coordination while pointing to perturbed targets, *Journal of Motor Behavior* **26**: 135–146.

Carson, R. (1995). The dynamics of isometric bimanual coordination, *Experimental Brain Research* **105**: 465–476.

Chua, R. & Elliott, D. (1993). Visual regulation of manual aiming, *Human Movement Science* **12**: 365–401.

Chua, R. & Elliott, D. (1997). Visual control of target-directed movements, *Behavioral and Brain Sciences* **20**: 304–305.

Crossman, E. & Goodeve, P. (1983). Feedback control of hand-movements and Fitts' law, *Quarterly Journal of Experimental Psychology* **35A**: 251–278. Original work presented at the meeting of the Experimental Psychology Society, Oxford, England, July 1963.

Diedrich, F. J. & Warren, W. H. (1998). The dynamics of gait transitions: Effects of grade and load., *Journal of Motor Behavior* **30**(1): 60–78.

Eldred, E., Granit, R. & Merton, P. (1953). Supraspinal control of the muscle spindles and its significance, *Journal of Physiology* **122**: 498–523.

Elliott, D. & Allard, F. (1985). The utilization of visual feedback information during rapid pointing movement, *The Quarterly Journal of Experimental Psychology* **37A**: 407–425.

Elliott, D., Binsted, G. & Heath, M. (1999). The control of goal-directed limb movements: Correcting errors in the trajectory, *Human Movement Science* **18**: 121–136.

Elliott, D., Heath, M., Binsted, G., Ricker, K., Roy, E. & Chua, R. (1999). Goal directed aiming: Correcting force specification error with the right and left hands, *Journal of Motor Behavior* **31**(4): 309–324.

Feldman, A. (1966a). Functional tuning of the nervous system with control of movements or maintenance of steady posture. II. controllable parameters of the muscle, *Biophysics* **11**: 565–578.

Feldman, A. (1966b). Functional tuning of the nervous system with control of movements or maintenance of steady posture. III. mechanographic analysis of execution by man of the simplest motor task, *Biophysics* **11**: 667–675.

Feldman, A. (1986). Once more on the equilibrium-point hypothesis (λ-model) for motor control, *Journal of Motor Behavior* **18**: 17–54.

Fitts, P. (1954). The information capacity of the human motor system in controlling amplitude of movement, *Journal of Experimental Psychology* **47**: 381–391.

Fitts, P. & Peterson, J. (1964). Information capacity of discrete motor responses, *Journal of Experimental Psychology* **67**: 103–112.

Flash, T. & Hogan, N. (1985). The coordination of arm movements: An experimentally confirmed mathematical model, *Journal of Neuroscience* **5**: 1688–1703.

Franz, E., Ivry, R. & Helmuth, L. (1996). Reduced timing variability in patients with unilateral cerebellar lesions during bimanual movements, *Journal of Cognitive Neuroscience* **8**: 107–118.

Fuchs, A. & Kelso, J. (1994). A theoretical note on models of interlimb coordination, *Journal of Experimental Psychology: Human Perception and Performance* **20**: 1088–1097.

Gaudiano, P. & Grossberg, S. (1991). Vector associative maps: Unsupervised real-time error-based learning and control of trajectories, *Neural Networks* **4**: 147–183.

Georgopolous, A., Kalaska, J., Caminiti, R. & Massey, J. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex, *Journal of Neuroscience* **2**: 1527–1537.

Georgopolous, A., Schwartz, A. & Kettner, R. (1986). Neuronal population coding of movement direction, *Science* **260**: 1416–1419.

Ghahramani, Z. (2000). An introduction to hidden Markov models and Bayesian networks, *International Journal of Pattern Recognition and Artificial Intelligence* **15**(1): 9–42.

Ghahramani, Z. & Jordan, M. (1997). Factorial hidden Markov models, *Machine Learning* **29**: 245–273.

Gibson, J. (1966). *The senses considered as perceptual systems*, Houghton Mifflin, Boston, MA.

Gibson, J. (1979). *The ecological approach to visual perception*, Houghton Mifflin, Boston, MA.

Grossberg, S., Roberts, K., Aguilar, M. & Bullock, D. (1997). A neural model of multimodal adaptive saccadic eye movement control by superior colliculus, *Journal of Neuroscience* **17**: 9706–9725.

Haken, H. (1977). *Synergetics*, Springer-Verlag, Heidelberg.

Haken, H. (1983). *Advanced Synergetics*, Springer-Verlag, Heidelberg.

Haken, H., Kelso, J. & Bunz, H. (1985). A theoretical model of phase transitions in human hand movements, *Biological Cybernetics* **51**: 347–356.

Henry, F. & Rogers, D. (1960). Increased response latency for complicated movements and a "memory drum" theory of neuromotor reaction, *Research Quarterly* **31**: 448–458.

Hewitt, C. (1985a). The challenge of open systems, *BYTE* **10**(4): 223–242.

Hewitt, C. (1985b). Offices are open systems, *ACM Transactions on Office Information Systems* **4**(3): 271–287.

Hick, W. (1952). On the rate of gain of information, *Quarterly Journal of Experimental Psychology* **4**: 11–26.

Hyman, R. (1953). Stimulus information as a determinant of reaction time, *Journal of Experimental Psychology* **45**: 188–196.

Ivry, R. & Keele, S. (1989). Timing function of the cerebellum, *Journal of Cognitive Neuroscience* **1**: 136–151.

Ivry, R., Keele, S. & Diener, H. (1988). Dissociation of the lateral and medial cerebellum in movement timing and movement execution, *Experimental Brain Research* **73**(1): 167–180.

Jankovic, J. & Tolosa, E. (1993). *Parkinson's disease and movement disorders*, 2nd edn, Williams and Wilkins, Baltimore, MD.

Jong, W. P. D. & Galen, G. P. V. (1997). Are speed/accuracy trade-offs caused by neuromotor noise , or not?, *Behavioral and Brain Sciences* **20**: 306–307.

Jordan, M. (1996). Attractor dynamics and parallelism in a connectionist sequential machine, *Proceedings of the Eighth Annual Conference of the Cognitive Society* pp. 531–546.

Kandel, E. R., Schwartz, J. H. & Jessell, T. M. (1991). *Principles of Neural Science*, third edn, Appleton and Lange, Conneticut.

Kawato, M. (1990). Computational schemes and neural networks for formation and control of multijoint arm trajectory, *in* T. Miller, R. Sutton & P. Werbos (eds), *Neural Networks for Control*, MIT Press, Cambridge, MA, pp. 197–228.

Kawato, M. (1996). Trajectory formation in arm movements: Minimization principles and procedures, *in* H. N. Zelaznik (ed.), *Advances in Motor Control and Learning*, Human Kinetics Publishers, Champaign, IL, chapter 9, pp. 225–260.

Keele, S. (1968). Movement control in skilled motor performance, *Psychological Bulletin* **70**: 387–403.

Kelso, J. (1994). The informational character of self-organized coordination dynamics, *Human Movement Science* **13**: 415–440.

Kelso, J. A. S. (1984). Phase transitions and critical behavior in human bimanual coordination, *American Journal of Physiology: Regulatory Integrative Comparative Physiology* **15**: R1000 – R1004.

Kugler, P., Kelso, J. & Turvey, M. (1980). On the concept of coordinative structures as dissipative structures: I. Theoretical lines of convergence, *in* G. Stelmach & J. Requin (eds), *Tutorials in motor behavior*, North Holland, Amsterdam, pp. 3–47.

Kugler, P., Kelso, J. & Turvey, M. (1982). On the control and coordination of naturally developing systems, *in* J. Kelso & J. Clark (eds), *The development of movement control and coordination*, Wiley, New York, NY, pp. 5–87.

Kugler, P. & Turvey, M. (1987). *Information, natural law, and the self-assembly of rhythmic movement*, Lawrence Erlbaum Associates, Hillsdale, NJ.

Latash, M. (1993). *Control of Human Movement*, Human Kinetics, Champaign, IL.

Latash, M. & Gottlieb, G. (1990). Compliant characteristics of single joints: Preservation of equifinality with phasic reactions, *Biological Cybernetics* **62**: 331–336.

Lee, D. & Reddish, D. (1981). Plummetin gannets: A paradigm of ecological optics, *Nature* **293**: 293–294.

Livingston & State, A. (1997). Magnetic tracker calibration for improving augmented reality registration, *Presence: Teleoperations & Virtual Environments* **6**(5): 532–546.

Mardia, K. (1972). *Statistics of directional data*, Academic Press, London, UK.

Marti, U. & Bunke, H. (2001). Using a statistical language model to improve performance of an HMM-based cursive handwriting recogntion system, *International Journal of Pattern Recognition and Artificial Intelligence* **15**(1): 66–91.

Massone, L. & Bizzi, E. (1989). A neural network model for limb trajectory formation, *Biological Cybernetics* **61**: 417–425.

McMahon, T. (1984). *Muscles, reflexes,and locomotion*, Princeton University Press, Princeton, NJ.

Merton, P. (1953). Speculations on the servo-control of movements, *in* J. Gray & G. Wolstenholm (eds), *The spinal cord*, Little Brown, Boston, MA, pp. 183–198.

Meyer, D., Abrams, R., Kornblum, S., Wright, C. & Smith, J. (1988). Optimality in human motor performance: Ideal control of rapid aimed movements, *Psychologica Review* **95**(3): 340–370.

Meyer, D., Smith, J. & Wright, C. (1982). Models for the speed and accuracy of aimed movements, *Psychological Review* **89**(5): 449–482.

Michaels, C. & Beek, P. (1995). The state of ecological psychology, *Ecological Psychology* **7**(4): 259–278.

Mitchell, T. M. (1997). *Machine Learning*, McGraw-Hill Series in Computer Science, McGraw-Hill, Boston, MA.

Mozer, M. (1995). A focused BACKPROPOGATION algorithm for temporal pattern recognition, *in* Y. Chauvin & D. Rumelhart (eds), *Backpropogation,: Theory, architectures, and applications*, Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 137–169.

Narendra, K. S. (1995). Adaptive control: neural network applications, *in* M. A. Arbib (ed.), *The Handbook of Brain Theory and Neural Networks*, MIT Press, Cambridge,MA, pp. 69–73.

Newell, K. (2001). Where have we been and where are we going? developing a new theory of motor learning for 2001 and beyond, Panel Discussion. Annual meeting of the North American Society for the Psychology of Sport and Physical Activity.

Park, H., Sin, B., Moon, J. & Lee, S. (2001). A 2-D HMM method for offline handwritten character recognition, *International Journal of Pattern Recognition and Artificial Intelligence* **15**(1): 92–107.

Plamondon, R. (1995a). A kinematic theory of rapid human movements part i. Movement representation and generation, *Biological Cybernetics* **72**: 295–307.

Plamondon, R. (1995b). A kinematic theory of rapid human movements part ii. Movement time and control, *Biological Cybernetics* **72**: 309–320.

Plamondon, R. (1998). A kinematic theory of rapid human movements part iii. Kinetic outcomes, *Biological Cybernetics* **78**: 133–145.

Plamondon, R. & Alimi, A. (1997). Speed/accuracy trade-offs in target-directed movements, *Behavioral and Brain Sciences* **20**: 279–349.

Plamondon, R. & Privitera, C. (1995). A neural model for generating and learning a rapid movement sequence, *Bioloigcal Cybernetics* **74**: 117–130.

Polit, A. & Bizzi, E. (1978). Processes controlling arm movements in monkey, *Science* **201**: 1235–1237.

Raab, F. H., Blood, E. B., Steiner, T. O. & Jones, H. R. (1979). Magnetic position and orientation tracking systems, *IEEE Transactions on Aerospace and Electronics Systems* **15**(5): 709–717.

Rabiner, L. & Juang, B.-H. (1993). Theory and implementation of Hidden Markov Models, *Fundamentals of Speech Recognition*, Prentice Hall, New York, NY, chapter 6, pp. 321–389.

Raibert, M. (1977). Motor control and learning by the state space model, *Technical Report AI-TR-439*, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, Cambridge, MA.

Rimey, R. & Brown, C. (1990). Selective attention as sequential behavior: Modeling eye movement with an augmented hidden Markov model, *University of Rochester Computer Science Technical Report* **327**.

Rosenbaum, D., Engelbrecht, S. E., Brushe, M. M. & Loukopoulos, L. D. (1993). A model for reaching control, *Acta Pschologica* **82**: 237–250.

Rosenbaum, D., Meulenbroek, R. G., Vaughan, J. & Jansen, C. (1999). Coordination of reaching and grasping by capitilizing on obstacle avoidance and other constraints, *Experimental Brain Research* **128**: 92–100.

Schmidt, R. (1976). Control processes in motor skills, *Exercise and Sport Sciences Review* **4**: 229–261.

Schmidt, R. (1982). More on motor programs, *in* J. Kelso (ed.), *Human motor behavior: An introduction*, Erlbaum, Hillsdale, NJ, pp. 189–217.

Schmidt, R. (1985). The search for invariance in skilled movement behavior, *Research Quarterly for Exercise and Sport* **56**: 188–200.

Schmidt, R. A. (1988). *Motor control and learning*, Human Kinetics, Champaign, IL.

Schmidt, R. A., Zelaznik, H. N., Hawkins, B., Frank, J. S. & Quinn, J. T. (1979). Motor output variability: A theory for the accuracy of rapid motor acts, *Psychological Review* **86**: 415–451.

Schmidt, R. & Turvey, M. (1994). Phase-entrainment dynamics of visually coupled rhythmic movements, *Biological Cybernetics* **70**: 369–376.

Schmidt, R. & Turvey, M. (1995). Models of interlimb coordination - equilibria, local analyses, and spectral patterning: comment on Fuchs and Kelso (1994), *Journal of Experimental Psychology: Human Perception and Performance* **21**(2): 432–443.

Schmidt, R., Zelaznik, H. & Frank, J. (1978). Sources of inaccuracy in rapid movement, *in* G. Stelmach (ed.), *Information processing in motor control and learning*, Academic Press, New York, NY, pp. 183–203.

Schoner, G. (1990). A dynamic theory of coordination of discrete movement, *Biological Cybernetics* **63**: 257–270.

Schöner, G., Zanone, P. & Kelso, J. (1992). Learning as a change of coordination dynamics: Theory and experiment, *Journal of Motor Behavior* **24**: 29–48.

Schwartz, A., Kettner, R. & Georgopolous, A. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space, i: Relations between single cell discharge and direction of movement, *Journal of Neuroscience* **8**: 2913–2927.

Sherwood, D. (1986). Impulse characteristics in rapid movement: Implications for impulse-variability models, *Journal of Motor Behavior* **18**: 188–214.

Stacey, D. (1998). *Introduction to Neuro-Computing: Artificial Neural Networks*, http://www.calresco.org/tuturial.htm.

Sternad, D., Amazeen, E. & Turvey, M. (1996). Diffusive, synaptic, and synergetic coupling: An evaluation through inphase and antiphase, *Journal of Motor Behavior* **28**: 225–269.

Swinnen, S., Beirinckx, M., Meugens, P. & Walter, C. (1991). Dissociating structural and metrical specifications of bimanual movement, *Journal of Motor Behavior* **23**: 263–279.

Tuller, B. & Kelso, J. (1989). Environmentally-specified patterns of movement coordination in normal and split-brain subjects, *Experimental Brain Research* **75**: 306–316.

Uno, Y., Kawato, M. & Suzuki, R. (1989). Formation and control of optimal trajectory in human multijoint movement: Minimum torque-change model, *Biological Cybernetics* **61**: 89–101.

Veeneman, D. & Lee, T. (2000). Bimanual coordination of a confederate/coordinator unit. Personal communication.

Verschueren, S., Swinnen, S., Cordo, P. & Dounskaia, N. (1999). Proprioceptive control of multijoint movement: Bimanual circle drawing, *Experimental Brain Research* **127**: 182–192.

Wallace, S. (1996). Dynamic pattern perspective of rhythmics movement: An introduction, *in* H. N. Zelaznik (ed.), *Advances in Motor Control and Learning*, Human Kinetics Publishers, Champaign, IL, pp. 155–195.

Warren, W. H., Kay, B. A. & Yilmaz, E. H. (1996). Visual control of posture during walking: Functional specificity, *Journal of Experimental Psychology: Human Perception and Performance* **22**(4): 818–838.

Wilson, A. & Bobick, A. (2001). Hidden Markov models for modeling and recognizing gesture under variation, *International Journal of Pattern Recognition and Artificial Intelligence* **15**(1): 92–107.

Wing, A. (1973). Response delays and the timing of discrete motor responses, *Perception and Psychophysics* **14**: 5–12.

Woodworth, R. S. (1899). The accuracy of voluntary movements, *Psychological Review Monographs* **3**(13).

Wright, C. & Meyer, D. (1983). Conditions for a linear speed-accuracy trade-off in aimed movements, *Quarterly Journal of Experimental Psychology* **35**: 279–296.

Zaal, F. T., Bootsma, R. J. & van Wieringen, P. C. (1998). Coordination in prehension: Information-based coupling of reaching and grasping, *Experimental Brain Research* **119**: 427–435.

Zaal, F. T., Bootsma, R. J. & van Wieringen, P. C. (1999). Dynamics of reaching for stationary and moving objects: Data and model, *Journal of Experimental Psychology: Human Perception and Performance* **25**(1): 149–161.

Appendix A

Polhemus Operations

A.1 System Configuration

The Polhemus Ultratrak Pro* six dimensional movement sensing system is employed for a wide variety of tasks ranging from animation and computer graphics to medical applications. In order to capture movement, three orthogonally arranged transmitter coils (magnetic-dipole field source) generate a pulsed field. Within the field, the position and orientation of a set of receivers (up to 32), also composed of three coils, are sampled (Raab, Blood, Steiner & Jones 1979). Sampling rates of either 60 or 120 Hz. are possible depending upon the number of active sensors and the system configuration.

Magnetic system have several key features that make them an attractive alternative to optical systems (e.g. Optotrak†). By collecting 6-degree of freedom data (x, y, z, yaw, pitch, roll) collection, magnetic motion tracking allows a more rich description of movement. The strongest benefit is occlusion free recording; optical systems are reliant upon line-of site between the camera(s) and the transmitter, thus reducing the freedom

*http://www.polhemus.com/; Hereafter "Polhemus"
†http://www.ndigital.com

of the subject to move naturally. One drawback however, is a high susceptibility to field distortions due to ferrous metals, thus lowering the accuracy and precision of measurements; these errors may be large depending upon the size and proximity of the iron to the receiver coils. Errors have been substantially reduced however, by calibration of the collection area, with precision in the order of ± 1 mm and $\pm 2^o$ for position and orientation respectively within 2 m of source coil (Livingston & State 1997, Boulic 1999).

A.2 System Calibration

As mentioned in the previous section, anisotropic distortions in the magnetic field of sensing systems such as Polhemus can cause substantial measurement error. Recently, Bryson (1992) investigated calibration for a Polhemus system using two methods.

The first, the least squares polynomial model, may be expressed as:

$$p_x' = poly_x(p_x, p_y, p_z) \tag{A.1}$$

$$p_y' = poly_y(p_x, p_y, p_z) \tag{A.2}$$

$$p_z' = poly_z(p_x, p_y, p_z) \tag{A.3}$$

where p_x, p_y, p_z correspond to position coordinates. After testing polynomials ranging in order from 1 to 8, a 4^{th} order solution produced the most accurate result. The alternative is table lookup; this method is currently employed by the Polhemus corporation for calibration of the Ultratrak Pro. During table lookup an offset value is assigned to each incoming point based upon the weighted sums of errors (offsets) observed at known fiducial points (Equation A.4):

$$\delta\tilde{x}_i = \sum_{j \in N_i} \phi(d_{ij}) \delta\tilde{x}_j \tag{A.4}$$

where $\phi(d_{ij})$ defines the function of distance between fiducial position $\tilde{X}_j$ in a neighbourhood of i, N_i, and a test position $\tilde{X}_i$. The final term, $\delta\tilde{X}_j$, being the known adjustment for fiducial $\tilde{X}_j$. Due to the reliance of such a model on the distance between fiducials for accurate interpolation, a dense calibration grid would be advantageous; the calibration parameters supplied by Polhemus however entail grid dimensions of $1ft^3$.

A.2.1 Gaussian Mixtures

The initial calibration method applied to data collection incorporated the non-linearities of the polynomial regression and the local-list nature of the lookup table interpolation into a single Gaussian Field Mixture Model (GFM). In this method, field distortions were modeled as tri-variate normal probability density functions (Equation A.5) surrounding fiducials points.

$$p_{ij} = \frac{1}{(|\Gamma|2\pi)^{3/2}} e^{-(\tilde{X}_j - \tilde{X}_i)\Gamma^{-1}(\tilde{X}_j - \tilde{X}_i)^T} \tag{A.5}$$

Here, $\tilde{X}_i$ corresponds to the position of fiducial point i and $\tilde{X}_j - \tilde{X}_i$ to the difference vector between data point j and fiducial i. $\tilde{X}^T$ is the transpose of vector $\tilde{X}$ with Γ corresponding to the 3D covariance matrix which defines shape of the Gaussian ($\delta\tilde{X}$).

Using a harmonic mean, the net effect of the error distributions associated with all neighbouring fiducials ($\delta\tilde{X}_j$) may be combined into a correction offset to be applied to the incoming data point (Equation A.6, see Figure A.1).

$$\delta\tilde{X}_j = \frac{\sum_i p_{ij}\delta\tilde{X}_i}{\sum_i p_{ij}} \tag{A.6}$$

The aim of this calibration procedure was to minimize difference between the veridical location of the sensor and its updated value. The parameters of the covariance matrices

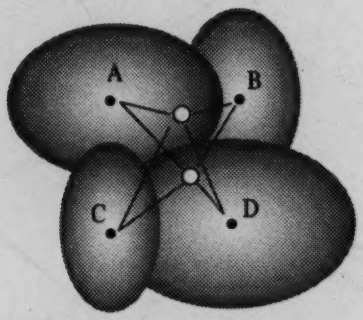

Figure A.1: As a new data point arrives, the harmonic mean of the offsets being applied to the surrounding fiducials is applied to that point (Caelli et al. 2000).

were independently modified to further rotate and scale each distribution to maximize fit to the data (Caelli et al. in submission).

However, while this mixture method does indeed demonstrate high spatial resolution (< 0.5 cm) it induces substantial computational load. In fact, the time required to locate a point within the field, estimate its nearest distortion distributions, and generate a corrected value served to dissable a Pentium (600 MHz, 126 mgb RAM) PC when more than 5 sensors were in use. Do to this fatal shortfall, other methods were investigated.

A.3 A Neural Network Calibration Algorithm

Since feed-forward neural networks are well known locally adaptive polynomial fitting procedures, a calibration procedure based on such an architecture was implemented (Figure A.2). A step-wise approach was implemented where different calibrations were performed using an increasing number of hidden units until best performance was obtained at the cross-over between RMS error on training and test data (Mitchell 1997).

Thirty-six (36) hidden units were found to cover the variations in the magnetic field, due presumably to the distribution of metal in the environment. For an volume inscribed by

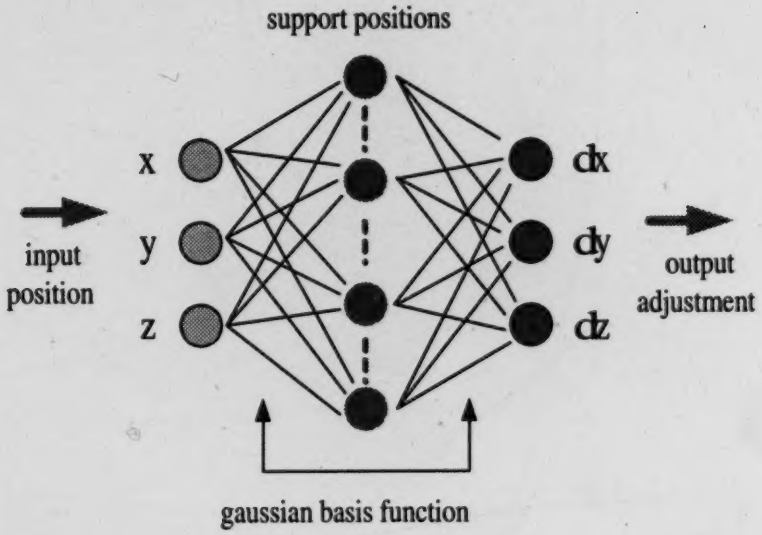

Figure A.2: Neural network calibration model. Here the hidden units are used to as local polynomial interpolation functions.

1.5 m from the transmitter, this resulted in a least overall error of ±5mm. Unlike the GFM approach, the majority of computation was completed during model estimation; at runtime, this approach generated very little processing delay. Returning to earlier methods of calibration, note that although this approach implements a neural network, the effective behaviour of calibration is very similar to that of the polynomial estimates of Bryson (1992).

Appendix B

Example gDBN

Included here is an example gDBN parameter set. As is evident by the weighting matrix, this model was trained with a 10%-10% connection strength. The data used for model update was an IPpost file (i.e. post transition in-phase data).

[Model 1]
M= 18.00
N= 18.00
A:

0.45	0.47	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.01	0.01	0.00	0.01	0.01	0.00	0.01	0.01
0.02	0.02	0.84	0.00	0.00	0.00	0.01	0.01	0.01	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01
0.02	0.01	0.01	0.70	0.15	0.00	0.01	0.01	0.01	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01
0.02	0.01	0.01	0.01	0.64	0.19	0.01	0.01	0.01	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.02
0.02	0.01	0.01	0.00	0.01	0.68	0.15	0.01	0.01	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.02
0.02	0.01	0.01	0.00	0.01	0.01	0.80	0.03	0.01	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.02
0.02	0.01	0.01	0.00	0.01	0.00	0.01	0.83	0.02	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01
0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.60	0.03	0.17	0.06	0.01	0.01	0.01	0.01	0.01	0.02
0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.07	0.65	0.10	0.02	0.01	0.01	0.01	0.01	0.01	0.02
0.03	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.03	0.08	0.69	0.02	0.01	0.01	0.01	0.01	0.01	0.02
0.02	0.01	0.01	0.00	0.01	0.00	0.01	0.01	0.01	0.01	0.11	0.75	0.01	0.01	0.01	0.01	0.01	0.01
0.02	0.01	0.01	0.00	0.01	0.00	0.01	0.01	0.01	0.01	0.02	0.02	0.82	0.01	0.01	0.01	0.01	0.02
0.02	0.01	0.01	0.00	0.01	0.00	0.01	0.01	0.01	0.01	0.02	0.01	0.01	0.65	0.19	0.01	0.01	0.02
0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.01	0.03	0.01	0.01	0.01	0.38	0.42	0.01	0.02
0.03	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.01	0.03	0.02	0.01	0.01	0.01	0.57	0.20	0.02
0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.70	0.13
0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.83
0.79	0.01	0.01	0.00	0.00	0.00	0.00	0.01	0.01	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.09

B:

0.77	0.23	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.95	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.75	0.25	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.66	0.34	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.43	0.57	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.66	0.34	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.76	0.24	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00

pi:

0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06

[Model 2]
M= 18.00
N= 18.00
A:

0.47	0.49	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.01	0.01	0.80	0.13	0.00	0.00	0.01	0.00	0.00	0.01	0.00	0.01	0.00	0.01	0.00	0.00	0.00	0.00
0.01	0.01	0.00	0.60	0.32	0.00	0.01	0.00	0.00	0.01	0.00	0.01	0.00	0.01	0.00	0.00	0.00	0.01
0.01	0.01	0.00	0.00	0.41	0.49	0.01	0.01	0.00	0.01	0.01	0.01	0.00	0.01	0.00	0.00	0.00	0.01
0.02	0.01	0.01	0.01	0.01	0.63	0.26	0.01	0.01	0.01	0.01	0.01	0.00	0.01	0.01	0.00	0.00	0.01
0.01	0.01	0.00	0.00	0.01	0.00	0.87	0.05	0.00	0.01	0.01	0.01	0.00	0.01	0.00	0.00	0.00	0.01
0.01	0.01	0.00	0.00	0.00	0.00	0.01	0.88	0.04	0.01	0.00	0.01	0.00	0.01	0.00	0.00	0.00	0.01
0.01	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.92	0.01	0.00	0.01	0.00	0.01	0.00	0.00	0.00	0.01
0.01	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.13	0.81	0.01	0.00	0.00	0.01	0.00	0.00	0.00	0.00
0.01	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.32	0.63	0.01	0.00	0.00	0.00	0.00	0.00	0.00
0.01	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.01	0.01	0.92	0.00	0.01	0.00	0.00	0.00	0.00
0.01	0.01	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.01	0.00	0.01	0.88	0.04	0.00	0.00	0.00	0.01
0.01	0.01	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.01	0.00	0.01	0.00	0.81	0.11	0.00	0.00	0.01
0.01	0.01	0.00	0.00	0.01	0.00	0.01	0.00	0.00	0.01	0.00	0.01	0.00	0.01	0.29	0.63	0.00	0.01
0.02	0.01	0.01	0.01	0.01	0.00	0.02	0.01	0.01	0.01	0.01	0.01	0.00	0.01	0.04	0.01	0.80	0.01
0.02	0.01	0.01	0.00	0.01	0.00	0.01	0.01	0.01	0.01	0.01	0.01	0.00	0.01	0.01	0.00	0.71	0.17
0.02	0.01	0.00	0.00	0.01	0.00	0.01	0.00	0.00	0.01	0.00	0.01	0.00	0.01	0.00	0.00	0.00	0.91
0.91	0.02	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.01	0.00	0.01	0.00	0.01	0.00	0.00	0.00	0.01

B:

0.89	0.11	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.02	0.98	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.05	0.95	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.10	0.49	0.40	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.05	0.95	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.77	0.23	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00
0.07	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.93

pi:

0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06

[Model 1on2]
A:

0.51	0.37	0.12	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.27	0.43	0.30	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.19	0.35	0.30	0.15	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.06	0.37	0.38	0.19	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.06	0.25	0.63	0.06	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.12	0.19	0.62	0.06	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.16	0.22	0.63	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.15	0.33	0.47	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.34	0.66	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.57	0.35	0.08	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.63	0.32	0.05	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.67	0.27	0.06	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.60	0.34	0.05	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.58	0.29	0.06	0.06	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.33	0.42	0.25	0.00
0.04	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.38	0.44	0.15
0.02	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.35	0.63
0.66	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.34

[Model 2on1]
A:

0.39	0.43	0.18	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.21	0.45	0.34	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.29	0.41	0.18	0.12	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.20	0.65	0.14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.23	0.38	0.33	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.51	0.40	0.10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.50	0.29	0.22	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.56	0.26	0.18	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.38	0.11	0.52	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.22	0.36	0.42	0.00	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.30	0.70	0.00	0.00	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.05	0.25	0.60	0.10	0.00	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.06	0.12	0.53	0.30	0.00	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.06	0.06	0.47	0.41	0.00
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.19	0.71	0.10
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.08	0.08	0.83
0.35	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.12	0.53
0.85	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.15

[Weightings]
C:

1.00	0.10
0.10	1.00